INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 20 Motion Control, Robotics, and Autonomous Systems

CHAPTER 8

A Low Earth Orbit Simulator Flown on Mean Elements and Collision Probability

Pages 187–210

Edited by John D. Hedengren

How to cite this chapter. "A Low Earth Orbit Simulator Flown on Mean Elements and Collision Probability," in *Motion Control, Robotics, and Autonomous Systems*, John D. Hedengren, Ed., vol. 20 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 8, pp. 187–210.

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

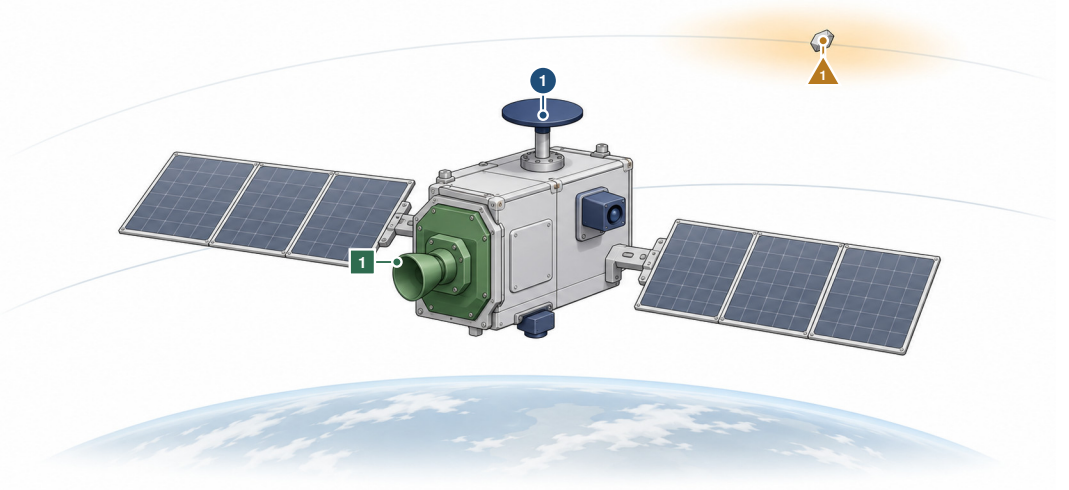
The app for this chapter
https://apmonitor.com/apps/leo-orbit-control/

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

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

CHAPTER 8

A Low Earth Orbit Simulator Flown on Mean Elements and Collision Probability



The controller moves 1 electric (Hall) thruster that fires every burn

The controller reads 1 satellite navigation antenna for the position fixes

What disturbs it 1 debris fragment, its position known only within the amber haze

Notice the fragment far from the satellite: whether its pass is safe depends on how uncertain its position is, not on the distance alone. Illustration, not to scale.

IN PLAIN TERMS Keeping a Satellite Up and Out of Harm's Way

The main idea. Thin air at the top of the sky slowly drags a low satellite down, so a small engine gives it a gentle push now and then. Its height is judged by the average over one lap, since the bulge of the Earth makes it swing on every lap.

The challenge. Bits of old rockets cross its path, and nobody knows exactly where they are. A piece that looks far away can still be a danger, so each push is planned hours ahead on the odds of a hit.

Why it matters. The fuel is small and every crash makes more debris, so each push must count.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the two history segments of Section 8.2, the try-it segment of scenario 1 (Section 8.4), the fact on the miss distance (Section 8.5), the try-it segment of scenario 3 and the fact on the campaign (Section 8.7), and the key points
Study	50 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the measured cases (Section 8.7), Technology in Practice (Section 8.8), Limitations (Section 8.9) and the open challenge

The control question A satellite in low Earth orbit is flown on a mean altitude that no instrument reports, a drag that has to be forecast and the collision probability of a pass hours ahead. When must a burn be decided, how large must it be, and on which prediction should it be judged?

Learning objectives

- Explain why orbit maintenance works on the one-orbit mean altitude and why one prograde burn raises the far side of the orbit.
- Compute the change of the mean, the burn time and the xenon of a reboost, and the drift that an along-track burn buys over a lead time.
- Evaluate the collision probability of a pass from its miss and its covariance, and show why a large miss distance can still be unsafe.
- Compare deadband and predictive maintenance, and nominal and uncertainty-aware avoidance, on the measured cases and the campaign.

Prerequisites Sections 1.4.6 (Kalman filter) and 1.4.10 (bound or commit), with Sections 1.6.1 and 1.6.4 for the runs. Chapter 7 helps but is not required.

The app https://apmonitor.com/apps/leo-orbit-control/

See also Volume 5, Chapter 13 for a vehicle steered on a filter's estimate of a position nothing aboard measures. Volume 9, Chapter 7 for a decision taken days ahead on an uncertain forecast under a chance constraint, and Volume 9, Chapter 5 for releases made one travel time ahead of the demand.

8.1 Introduction

A feedback loop compares a measurement with a setpoint. A satellite in low Earth orbit is flown instead on three quantities that no instrument reports. The first is the mean altitude, an average over the last revolution, because the oblateness of the Earth adds periodic terms to every orbital element and the motion of a close satellite is written in mean elements with those terms added [1]. The second is the drag that lowers the mean, set by an upper-atmosphere density that the first satellites were used to measure from the contraction of their own orbits [2] and that density models still have to supply [3]. The third is the collision probability of a debris pass announced hours ahead, computed from two uncertain orbits in the plane normal to their relative velocity [4, 5].

The simulated satellite is a 120 kg Earth-observation spacecraft in a sun-synchronous orbit at a mean altitude of 500 km and an inclination of 97.4°, pushed by a 20 mN Hall thruster that fires for a minute to change its velocity by 1 cm/s. Three properties separate it from a loop with a delay. With nothing but gravity acting, its altitude swings between 498.42 and 501.58 km twice an orbit while the mean holds within 1 m. A burn changes the altitude only on the far side of the orbit, half a revolution later, and moves the satellite along its track by about three times the velocity change for every second that follows [6]. The criterion for a safe pass is a probability, and a predicted miss distance does not measure it.

The simulator addresses two beliefs. One is that the altitude a satellite reports is the altitude it controls, and the other is that a pass predicted to miss by a wide margin is safe. The second has a history in operations, since Space Shuttle avoidance was commanded whenever a conjunction fell inside a box around the vehicle, while the International Space Station was designed to maneuver on a probability of collision [4]. Data of lower quality can even make a pass look safer, the effect called probability dilution [7]. The class of problem is general. A decision must satisfy a constraint with a stated probability, on a model whose uncertainty is carried forward to the moment the constraint applies, which is chance-constrained predictive control [8]. It bounds the unknown and commits to a prediction, two of the four ways of Section 1.4.10. A reservoir released against an inflow forecast, as in Volume 9, Chapter 7, and a power system committed against a load forecast have the same shape. A transfer function with a dead time represents none of it, because the controlled variables here are a running average and a probability.

The operational tools behind these decisions run on real catalogs and tracking data held by agencies and operators [9], and none of them is a plant on which a student can make the decisions and watch their consequences. This chapter casts orbit maintenance and collision avoidance as control problems whose controlled variables are predictions, in a browser simulator of a low Earth orbit satellite whose orbit filter sees only navigation fixes and whose catalog knows only its own tracking. The oblateness and drag models, the extended Kalman filter, the short-encounter collision probability and burn sizing on the linear relative motion are standard. Around them the simulator keeps a truth model apart from the operator's side, designs synthetic debris against the true trajectory with declared covariances that tracking shrinks, carries in the satellite's covariance the drag forecast and the execution errors of burns still to fly, and gives two planners that screen every burn against every object. Every setting the runs need is given, so that they can be reproduced. Eight cases and a four-day campaign meet forecasting problems and feedback problems one at a time. The debris is synthetic and the atmosphere illustrative, so nothing in the app is a live orbit forecast.

Section 8.2 reviews how drag, mean orbits and conjunction assessment came to be handled as they are. Section 8.3 gives the truth model of the satellite, the atmosphere and the synthetic debris, and Section 8.4 derives four consequences of that model for a planner: the oscillation, the decay, the response to a burn and the geometry of a pass. Section 8.5 gives the orbit filter, the screening and the collision probability, and Section 8.6 the maintenance and avoidance planners and the lower layer that flies their burns. Section 8.7 reports the eight cases and the campaign, Section 8.8 names the devices of real satellites, and Section 8.9 states the limits of the model.

8.2 Drag, Mean Orbits and Conjunction Assessment

Orbit decay turned the first satellites into instruments of the atmosphere they flew through. The contraction of the early orbits gave the density of the upper atmosphere at their heights and its variation from 1957 to 1961 [2]. The theory of a close satellite in the oblate field of the Earth was worked out in the same years, with the perturbations of the six elements expressed as functions of mean orbital elements and time, periodic terms of the first order and secular terms to the second [1]. The secular turning of the node is what makes a sun-synchronous orbit possible, and the split between periodic and secular motion is why maintenance is flown on mean elements. Methods of orbit determination and prediction, including the force models and the filters used here, are collected in [10].

> ⧖ HISTORY **Hall Thrusters Go to Orbit (1971 to 2019)**
>
> The first pair of Hall thrusters flew on a Soviet spacecraft in 1971, and more than 300 others flew between 1971 and 2018. The first built outside Russia flew in 2003, and the first built in the United States in 2006. Launches grew sharply in 2019, when two large constellations began to carry a Hall thruster on every satellite. The reasons given are a simple design, a reliability proven in flight, good efficiency, a high specific impulse and a high thrust for the power drawn. Every Hall thruster flown had run on xenon, the propellant of this chapter's satellite, until one of those constellations turned to krypton, which is far cheaper in the quantities that thousands of satellites need [11].

Density stayed the weak link. Thermospheric mass density, its measurement and its models are reviewed in [3]. In February 2022 a minor to moderate geomagnetic storm raised the density between 200 and 400 km by 50 to 125 %, and 38 of 49 newly launched satellites were lost, while pre-launch analyses with an empirical density model had underestimated the effect [12]. Station keeping has also moved on board. One study of a remote-sensing spacecraft at low altitude used a low-power Hall thruster for continuous in-track control, an extended Kalman filter with satellite navigation among its measurements, and a Lyapunov-based orbit law [13], an arrangement close to the one simulated here except that the app plans discrete burns.

The debris problem was framed in 1978 as a growth process, in which each collision produces fragments that raise the probability of the next, so that a belt of debris could form unless launch constraints and operational procedures delayed it [14]. The first collision between satellites came on 10 February 2009, when the active Iridium 33 met the retired Cosmos 2251, and by 1 December 2009 the U.S. tracking system had catalogued 1632 fragments from it [15]. Avoidance moved from distances to probabilities over the same decades. The Space Shuttle maneuvered when a conjunction fell inside a box extending 5 km along track and 2 km radially on either side of its predicted position, and the space station was designed to maneuver on a probability computed in the conjunction plane [4]. The theory of short encounters behind that probability is collected in [5], and an operational service for the collision risks of a fleet of satellites processes tracking data and updates ephemerides, covariances, conjunction geometries and risk estimates [9].

The weak points of that chain are the ones the app exposes. Uncertainty in the atmospheric density is usually left out when a covariance is carried to the time of closest approach, and it can change the probability of collision significantly, as the effect of a geomagnetic storm illustrates [16]. A covariance that is too large dilutes the probability, so that worse data appear to lower the risk, a symptom analyzed as false confidence [7]. The burn follows from a linear relation between the impulse and the relative motion in the encounter plane, which turns its design into the search for the least velocity change that meets a miss distance or a probability [17]. The linear relative motion is the terminal-guidance model of 1960 [6].

> ⧖ HISTORY **A Catalog Rebuilt for Probabilities (1986 to 2009)**
>
> The prediction of close approaches between objects in orbit began in support of the Space Shuttle's return to flight after the Challenger accident of 1986, drawing on a military network that tracked objects in orbit for missile defense. Its catalog of element sets covers tracked objects larger than about 10 cm in low orbit, but the analytic theory behind it gives no usable covariance. To compute a probability of collision for the space station, a more accurate catalog was built on special perturbations theory, with the covariances that the computation needs. Screening against that catalog has since grown, especially after a collision of two satellites in 2009, to cover every operational satellite in orbit [18].

8.3 Satellite, Atmosphere and Debris

The truth model carries the satellite in Earth-centered inertial coordinates with seven states, its position $\mathbf{r}$, its velocity $\mathbf{v}$ and the xenon m_p left in the tank. The orbit, the satellite's radial, along-track and cross-track axes $(\hat{\mathbf{R}}, \hat{\mathbf{T}}, \hat{\mathbf{N}})$, the forces, the navigation fix, the crossing geometry of a pass and its encounter plane are shown in Figure 8.1. The engine's clock runs in hours, so every rate below is multiplied by 3600 s/h in the code.

Table 8.1 assigns each signal its role. The controlled quantities are not states. The mean altitude is an average of the state over the last orbit, and the collision probability and the miss distance are predictions made for each pass.

8.3.1 Equations of Motion

Gravity is the central term plus the J_2 term of the oblateness, drag acts against the velocity relative to an atmosphere that turns with the Earth, and the thruster pushes along a direction fixed in the satellite axes, as shown in Equation (8.1).

$$
\begin{aligned}
\ddot{\mathbf{r}} &= -\frac{\mu}{r^3}\mathbf{r} - \frac{3 J_2 \mu R_E^2}{2 r^5}
\begin{bmatrix} x(1-5z^2/r^2) \\ y(1-5z^2/r^2) \\ z(3-5z^2/r^2) \end{bmatrix}
+ \mathbf{a}_D + \mathbf{a}_T \\
\mathbf{a}_D &= -\tfrac{1}{2}\rho B |\mathbf{v}_r| \mathbf{v}_r, \qquad \mathbf{v}_r = \mathbf{v} - \boldsymbol{\omega}_E \times \mathbf{r} \\
\mathbf{a}_T &= \frac{\eta\, g_b F \ell}{m_{dry} + m_p} \left(\hat{u}_R \hat{\mathbf{R}} + \hat{u}_T \hat{\mathbf{T}} + \hat{u}_N \hat{\mathbf{N}} \right)
\end{aligned}
\tag{8.1}
$$

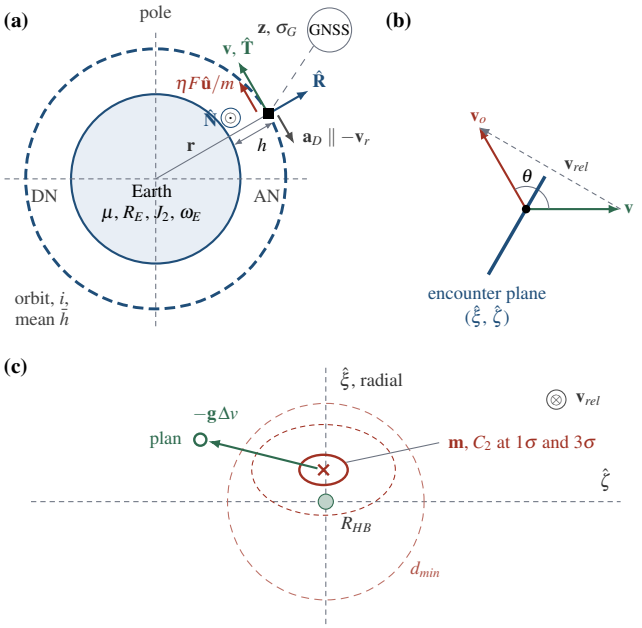


Figure 8.1 The simulated satellite, drawn out of scale. (a) The orbit plane with $\hat{\mathbf{N}}$ out of the page and the height exaggerated as in the app's orbit view, showing the equator through the ascending and descending nodes, the position $\mathbf{r}$ and altitude h, the satellite axes, the thrust along $\hat{\mathbf{u}}$, the drag against the wind relative to the turning atmosphere and the navigation fix $\mathbf{z}$. (b) A pass in the local horizontal plane, $\hat{\mathbf{R}}$ out of the page. The object crosses at the angle θ with the same speed, and the encounter plane is normal to the relative velocity. (c) The encounter plane at the time of closest approach (TCA), centered on the satellite and its combined hard-body circle, with the predicted miss $\mathbf{m}$ inside its combined covariance C_2, the miss-distance floor and the position a planned burn gives.

Table 8.1 Variables of the Satellite Plant by Role

Symbol	Role	Default	Range
$\bar{h}$	CV, true mean altitude	500.000 km	$\bar{h}_{SP} \pm \Delta h$
$\hat{\bar{h}}$	filter mean, planned on	—	—
P_c, d	CVs per pass	—	$\le P_{c,max}$, $\ge d_{min}$
$\bar{h}_{SP}$	band center	500 km	480–520 km
Δh	band half-width	250 m	0.05–2 km
u	MV in manual, thrust	0 %	±100 % of F
burns	MV in automatic	—	$\Delta v \le \Delta v_{max}$
s	DV, solar density scale	1	0–5
f_{st}	DV event, storm	1	peak 3
η	DV, thruster output	100 %	20–100 %
objects	DV, synthetic debris	per case	Table 8.6
$\mathbf{r}$, $\mathbf{v}$	states, inertial	mean circular	—
m_p	state, xenon	2.0 kg	0.2–10 kg
$\hat{\mathbf{x}}$	filter states	$\hat{k} = \hat{\eta} = 1$	clamped

where μ is the gravitational parameter, R_E the equatorial radius, ω_E the rotation of the Earth about its axis, $B = C_D A/(m_{dry} + m_p)$ the drag factor, F the rated thrust, $\ell \in [0,1]$ the share of the step the thruster fires, η its true output, g_b the magnitude factor of the burn and $\hat{\mathbf{u}}$ its direction with the pointing error included. The axes follow from the state, $\hat{\mathbf{R}} = \mathbf{r}/r$, $\hat{\mathbf{N}} = \mathbf{r} \times \mathbf{v}/|\mathbf{r} \times \mathbf{v}|$ and $\hat{\mathbf{T}} = \hat{\mathbf{N}} \times \hat{\mathbf{R}}$.

The thrust stops when the tank is empty. Integration is fixed-step fourth-order Runge-Kutta at 10 s. Halving the step moves the position after one day by 0.43 m, and an orbit under gravity alone keeps its energy to a relative 4.4×10^{-10} over a day.

8.3.2 Atmosphere and Storm

The density is a piecewise-exponential table after CIRA-72 for moderate solar activity, evaluated at the height above the ellipsoid with a geocentric latitude. It is scaled by the true solar-activity factor s and a storm profile, as shown in Equation (8.2).

$$\begin{aligned} \rho &= s\, f_{st}(t)\, \rho_j\, \mathrm{e}^{-(h_e - h_j)/H_j}, \qquad h_j \le h_e < h_{j+1} \\ h_e &= r - R_E \left(1 - f\, z^2/r^2\right) \end{aligned} \tag{8.2}$$

where f is the flattening and (h_j, ρ_j, H_j) the base height, base density and scale height of band j in Table 8.3. The storm multiplier rises linearly from 1 to its peak $p = 3$ over 2 h, holds for 12 h and returns over 6 h. The operator never sees s or f_{st}. Their only trace is in the orbit.

8.3.3 Thruster, Propellant and Battery

Planned burns fire the Hall thruster at its rated thrust, and in manual mode the operator's slider sets a share of it. The xenon flow follows the command, not the thrust delivered, and the battery pays for the power the thruster draws beyond the surplus of the solar arrays, as shown in Equation (8.3).

$$\dot{m}_p = -\frac{F\,\ell}{I_{sp}\, g_0}, \qquad \dot{E} = \begin{cases} -(P_{thr} - P_{sur})\ell & \text{firing} \\ P_{sur} & \text{idle} \end{cases} \tag{8.3}$$

with E clamped to $[0, E_{bat}]$. At 1400 s the flow is 1.46 mg/s, so each 1 cm/s of the 120 kg satellite costs 87.4 mg of xenon. A full battery allows 1.0 h of continuous firing, and the arrays sustain a duty cycle of 43 %. A thruster at $\eta = 0.55$ spends the same xenon and delivers 55 % of the change.

8.3.4 Mean Elements and the Initial State

The mean altitude and the mean eccentricity vector are running averages of the osculating radius and eccentricity vector over one period T_0 of the initial mean orbit, sampled every step, as shown in Equation (8.4).

$$\bar{h} = \frac{1}{L}\left(\sum_{j=0}^{N-1} |\mathbf{r}_{k-j}| + \phi\, |\mathbf{r}_{k-N}|\right) - R_E, \qquad L = \frac{T_0}{\Delta t} = N + \phi \tag{8.4}$$

where $\Delta t = 10$ s and the oldest sample carries the fractional weight ϕ. The mean eccentricity $\bar{e}$ is the magnitude of the averaged vector, and the mean perigee and apogee are $\bar{h} \mp (\bar{h} + R_E)\bar{e}$. The average removes the oscillation that J_2 puts into every element and lags half an orbit behind a burn. The initial state sits at the ascending node and is found by Newton's method on the same average, so that after one orbit of gravity alone the mean radius is $R_E + \bar{h}_0$ and the mean eccentricity vector is zero. The history before $t = 0$ is filled by coasting that state backward for one orbit. Starting instead on an osculating circle of radius $R_E + 500$ km would leave the mean altitude 4.87 km low and the mean eccentricity at 4.8×10^{-4}.

8.3.5 Synthetic Debris and Tracking

No real object or catalog is represented. Each object is created at its alert against the satellite's true trajectory, propagated from that moment with the truth model and the burns already committed. At the design TCA the object's velocity is the satellite's turned about $\hat{\mathbf{R}}$ by the crossing angle θ, so the relative velocity is horizontal, and its catalog position sits at the declared miss vector in the encounter plane. The true position differs from the catalog position by an offset δ drawn from the final catalog covariance, or from a tighter spread when the case declares one, in the object's own radial, along-track

Table 8.2 Satellite and Environment Parameters, All Shown and Adjustable in the Parameter Panel

Symbol	Meaning	Value	Range
m_{dry}, $m_{p,0}$	dry mass, xenon	118, 2.0 kg	50–500, 0.2–10
A, C_D	drag area, coefficient	2.4 m^2, 2.2	0.5–10, 1.5–3
F, I_{sp}	thrust, specific impulse	20 mN, 1400 s	5–100, 300–3000
P_{thr}, P_{sur}	power, array surplus	350, 150 W	100–1500, 20–1000
E_{bat}	battery for thrusting	200 Wh	20–2000
i	inclination	97.4°	0–180
$\bar{h}_0$	initial mean altitude	500 km	400–700
R_{HB}	combined hard-body radius	15 m	2–50
σ_G	fix noise per axis	5 m	0.5–50
σ_θ, σ_F	pointing, magnitude error	1°, 2 %	0–10, 0–20
t_{slew}	slew before a burn	5 min	0–30
σ_{drag}	drag forecast error	10 %	0–50
m_{alloc}	xenon allocated	30 g	1–1000

Table 8.3 Fixed Constants and the Density Bands, in the Source Only

$\mu = 3.986\,004\,418 \times 10^{14}\,m^3/s^2$	$R_E = 6378.137\,km$
$J_2 = 1.082\,626\,68 \times 10^{-3}$	$f = 1/298.257\,223\,563$
$\omega_E = 7.292\,115\,9 \times 10^{-5}\,rad/s$	$g_0 = 9.806\,65\,m/s^2$

h_j (km)	ρ_j (kg/m^3)	H_j (km)	h_j (km)	ρ_j (kg/m^3)	H_j (km)
300	2.418×10^{-11}	53.628	600	1.454×10^{-13}	71.835
350	9.518×10^{-12}	53.298	700	3.614×10^{-14}	88.667
400	3.725×10^{-12}	58.515	800	1.170×10^{-14}	124.64
450	1.585×10^{-12}	60.828	900	5.245×10^{-15}	181.05
500	6.967×10^{-13}	63.822	1000	3.019×10^{-15}	268.00

and cross-track axes. Both then move on two-body orbits from their states at TCA. The catalog gives the operator a mean offset $\hat{\delta}$ and a diagonal standard deviation σ per axis. Tracking updates, evenly spaced from the alert to a last update before TCA (one hour by default), shrink σ geometrically from σ_0 to σ_F through Kalman updates on a measurement of the true offset, as shown in Equation (8.5).

$$\begin{aligned} \sigma_k &= \sigma_0 \left(\sigma_F/\sigma_0\right)^{k/N_u}, \qquad r_k^2 = \left(\sigma_k^{-2} - \sigma_{k-1}^{-2}\right)^{-1} \\ \hat{\delta}_k &= \sigma_k^2 \left(\frac{\hat{\delta}_{k-1}}{\sigma_{k-1}^2} + \frac{\delta + r_k \varepsilon_k}{r_k^2} \right) \end{aligned} \tag{8.5}$$

where N_u counts the tracking updates and ε_k is a standard normal deviate. The operator sees $\hat{\delta}$ and σ and never δ. The declarations of every object are listed in Table 8.6.

8.3.6 Numbers

The parameters are listed in Table 8.2 and the constants and the density bands in Table 8.3. The dry mass and the xenon give $B = 0.044\,m^2/kg$, and 97.4° is the sun-synchronous inclination at 500 km. Every case starts from the mean circular state at $\bar{h}_0$ with the node at zero right ascension, a full battery, the tank at $m_{p,0}$ and the random streams seeded from one fixed seed, so every run replays bit for bit (Section 1.6.1).

8.3.7 Assumptions

Gravity stops at J_2. Higher harmonics, the Sun, the Moon and radiation pressure drive eccentricity and inclination over weeks and are absent, so multi-week eccentricity control and ground-track keeping

are outside the model.

The atmosphere is a table times two scales, illustrative rather than a forecast. Real density at these heights follows solar and geomagnetic activity [3], and the storm of February 2022 raised it by 50 to 125 % between 200 and 400 km [12], less than the tripling case 4 imposes. The height uses a geocentric latitude, within 25 m of the geodetic height, which moves the density by 0.04 %.

The satellite is a point mass with a constant drag area. Attitude, the slew itself and thermal limits are reduced to a fixed slew time and a pointing error drawn once per burn, so a burn never fails to start. The attitude of a small Earth-observation satellite is the subject of Chapter 7.

The debris is synthetic and moves on two-body orbits, with its own drag and J_2 folded into declared covariances. The design places every pass where the case wants it, so the geometry of a pass is chosen, not found by a catalog-wide screening.

Navigation uses white position fixes with no outages or clock errors, while an operational orbit determination system also processes velocity or range measurements with a richer force model [10].

8.4 The Orbit the Planners See

Four consequences of Equation (8.1) decide what a planner can do. The oscillation comes first. In case 1, with no air and no thrust, the altitude $|\mathbf{r}| - R_E$ swings between 498.42 and 501.58 km, the height above the ellipsoid between 501.6 and 519.5 km, and the osculating semimajor axis between 486.0 and 504.9 km above R_E, twice an orbit. The mean altitude stays within 0.98 m of 500 km for a day and the mean eccentricity at most 6.2×10^{-7}. The one-orbit average of the semimajor axis sits 4.56 km below the mean radius, so the app defines the mean on the radius, the quantity its altitude band refers to. The node turns at the rate of Equation (8.6).

$$\dot{\Omega} = -\tfrac{3}{2} n J_2 \left(\frac{R_E}{a} \right)^2 \cos i \tag{8.6}$$

A straight-line fit to the node over three days gives 0.9875° a day. Equation (8.6) gives 0.9854° with a at the mean radius and 0.9877° at the mean semimajor axis, within 0.02 % of the simulation, and the Sun appears to move 0.9856° a day.

> **TRY IT IN THE APP** **1 · Unforced Coast**
>
> Load the case, press Run and stay on the Orbit view. Watch the Altitude now trend swing up and down twice every orbit, between 498.42 and 501.58 km, while the Mean altitude trend stays flat. The chapter reports the mean within 0.98 m of 500 km for a day. Nothing pushes or drags this satellite, and the swing comes from the bulge of the Earth.

The second is the decay. Averaging the drag of Equation (8.1) over a near-circular orbit gives the fall of the mean semimajor axis in Equation (8.7), which is also the forecast the maintenance planner uses.

$$\frac{\mathrm{d}\bar{a}}{\mathrm{d}t} = -k\,\hat{\rho}_s\,\langle \rho \rangle \sqrt{\mu \bar{a}}\, B \left(1 - \frac{\omega_E \bar{a} \cos i}{v} \right)^2 \tag{8.7}$$

where $v = \sqrt{\mu/\bar{a}}$ and $\langle \rho \rangle$ is the table density averaged over 24 points around the orbit at the height above the ellipsoid. In the planner $k = \hat{k}$ and $\hat{\rho}_s$ is the operator's density scale, and with both at 1 the law gives the decay in the default atmosphere. At the tabulated density for exactly 500 km, $\rho \sqrt{\mu a} B$ gives 138.7 m a day. The orbit flies 1.6 to 19.5 km higher above the ellipsoid, where the air is thinner, and the average gives 120.4 m a day, while case 2 loses 115.4, 121.2 and 120.3 m on its first three days. The same law gives 270.9 and 630.3 m a day at 450 and 400 km.

The third is the burn. An along-track change Δv changes the mean semimajor axis and, for a burn made τ before a time of interest, the radial and along-track position at that time as the linear relative

motion [6] gives in Equation (8.8).

$$\Delta a = \frac{2a^2 v \Delta v}{\mu}, \qquad \begin{aligned} x &= \tfrac{2\Delta v}{n}\,(1-\cos n\tau) \\ y &= \tfrac{\Delta v}{n}\,(4\sin n\tau - 3n\tau) \end{aligned} \tag{8.8}$$

where $n = 1.1068 \times 10^{-3}\,\mathrm{rad/s}$ is the mean motion and the period is 94.62 min. At 500 km one m/s raises the mean 1.807 km. The point where a prograde burn fires stays where it was, and the far side rises by $2\Delta a$, so a single reboost makes the orbit eccentric. The secular term $-3\Delta v\,\tau$ is the along-track drift that avoidance buys, 10.8 m per hour for each mm/s.

🕮 **Pause.** Try Exercise 8.3 before reading on.

The fourth is the geometry of a pass. Two objects of equal speed crossing at the angle θ have a relative velocity that makes the angle $90^\circ - \theta/2$ with each along-track axis, so an along-track error enters the encounter plane multiplied by the factor of Equation (8.9).

$$\left|\hat{\mathbf{T}} \times \hat{\mathbf{v}}_{rel}\right| = \cos\frac{\theta}{2} \tag{8.9}$$

At 160° the factor is 0.174, so a nearly head-on pass hides most of an along-track error in the timing of the pass. At 100° it is 0.643 and the along-track error lands in the plane. The same projection applies to the along-track part of an avoidance burn, while its radial part enters whole.

8.5 Navigation, Screening and Collision Probability

8.5.1 Orbit Filter

The operator knows the orbit through an extended Kalman filter with eight states, $\hat{\mathbf{x}} = [\hat{\mathbf{r}}, \hat{\mathbf{v}}, \hat{k}, \hat{\eta}]$, where $\hat{k}$ scales the drag the operator's model predicts and $\hat{\eta}$ the thrust the lower layer commanded. The prediction integrates the operator's model with the same step, with J_2, drag at $\hat{k}\hat{\rho}_s$ times the table and the commanded thrust times $\hat{\eta}$. The covariance follows Equation (8.10).

$$\begin{aligned} &\Phi = I + A\Delta t + \tfrac{1}{2}(A\Delta t)^2, \qquad P \leftarrow \Phi P \Phi^{\top} + Q \\ &K = PH^{\top}(HPH^{\top} + \sigma_G^2 I)^{-1}, \qquad \hat{\mathbf{x}} \leftarrow \hat{\mathbf{x}} + K(\mathbf{z} - H\hat{\mathbf{x}}) \end{aligned} \tag{8.10}$$

where A holds the two-body gravity gradient $\mu r^{-3}(3\hat{\mathbf{R}}\hat{\mathbf{R}}^{\top} - I)$ and the sensitivities of the acceleration to $\hat{k}$ and $\hat{\eta}$, $H = [I_3\ 0]$ and $\mathbf{z} = \mathbf{r} + \varepsilon$ is a fix every 10 s with 5 m of white noise per axis. The covariance update uses the Joseph form, and $\hat{k}$ and $\hat{\eta}$ are clamped to $[0.2, 6]$ and $[0.2, 1.3]$ after each update. The process noise adds $q_a^2 \Delta t$ to each velocity variance, with $q_a = 2 \times 10^{-9}\,\mathrm{m/s^{1.5}}$, and while a burn fires another $60\Delta t\,(0.05 F\ell/\hat{m})^2$. It adds $q_k^2 \Delta t$ to the drag scale with $q_k = 1 \times 10^{-3}\,\mathrm{s^{-1/2}}$ and $q_\eta^2 \Delta t$ to the efficiency with $q_\eta = 1 \times 10^{-6}\,\mathrm{s^{-1/2}}$. The initial estimate is the truth plus random errors with standard deviations of 20 m and 0.02 m/s per axis, and the covariance starts diagonal at 20, 0.02, 0.3 and 0.1. The filter's own mean $\hat{\bar{h}}$ is the average of Equation (8.4) applied to $\hat{\mathbf{r}}$, and it is what the maintenance planner works on. Over hours 24 to 72 of case 2 the filter holds the position to 0.53 m and the velocity to 0.58 mm/s RMS while the fixes scatter 8.66 m RMS in three dimensions, and $\hat{k}$ stays between 0.92 and 1.07 against a true 1. The filter is the variance-chosen form of the fusion filters of Section 1.4.6, with the drag scale and the thrust efficiency as bias states. Volume 5, Chapter 13 steers an underwater vehicle on the same kind of filter, with fixes that arrive late.

8.5.2 Screening and the Time of Closest Approach

Every 30 min while a pass is pending, at every alert and tracking update and 15 min after every burn, the planner propagates the filter's estimate with the operator's model and the burns already committed. This operational trajectory runs to just past the last pass due within 30 h. For each catalogued object the closest approach is found on the 10 s grid within 2 min of the catalog TCA and refined on

straight-line relative motion. The encounter frame takes $\hat{\mathbf{e}}_z = \mathbf{v}_{rel}/|\mathbf{v}_{rel}|$, the satellite's radial direction projected into the plane as $\hat{\xi}$, and $\hat{\zeta} = \hat{\mathbf{e}}_z \times \hat{\xi}$, and the miss vector $\mathbf{m}$ is the object's position relative to the satellite resolved on those two axes. The last assessment 10 min before TCA is the one each pass is scored on.

8.5.3 The Satellite's Own Covariance at TCA

The filter covariance is rotated into the satellite's rotating axes and carried to TCA by the linear state transition $\Phi_{CW}(\tau)$ of the relative motion [6], and two terms are added, as shown in Equation (8.11).

$$C_{own} = \left[\Phi_{CW} P_H \Phi_{CW}^{\top}\right]_{rr} + \sigma_\kappa^2 \hat{a}_D^2 \, \mathbf{g}_a \mathbf{g}_a^{\top} + \sum_b G_b S_b G_b^{\top} \tag{8.11}$$

where $\mathbf{g}_a$ is the response at TCA to a unit along-track acceleration held for τ, $\hat{a}_D$ the modeled drag, and $\sigma_\kappa = (P_{kk}/\hat{k}^2 + \sigma_{drag}^2)^{1/2}$ the relative drag uncertainty of the filter and the forecast together. The sum, used only by the uncertainty-aware planner, carries the execution error of each burn still to fly through the impulse responses G_b, with S_b holding $(\sigma_F \Delta v)^2$ along the thrust, the filter's efficiency variance added, and $(\sigma_\theta \Delta v)^2$ across it. The drag term dominates early. At the alert of case 5, 9.7 h before TCA, it gives 407 m of along-track uncertainty at TCA, and at 4.2 h it gives 59 m. The combined covariance adds the catalog's diagonal covariance in the object's axes and is projected onto the plane, $C_2 = [\hat{\xi}\ \hat{\zeta}]^{\top}(C_{own} + C_{obj})[\hat{\xi}\ \hat{\zeta}]$.

8.5.4 Collision Probability in the Encounter Plane

The probability follows the short-encounter model [4, 5]. At 10.8 to 15.0 km/s the 30 m hard-body diameter is crossed in 2 to 3 ms. The relative motion during the encounter is therefore a straight line at the relative velocity found at TCA, the position errors of both objects are Gaussian and frozen while it lasts, velocity errors are not carried, and the two errors are independent so that their covariances add. Both objects are spheres whose radii add to R_{HB}, and the probability is the mass of a two-dimensional Gaussian inside a disc, as shown in Equation (8.12).

$$P_c = \iint_{|\mathbf{x}| \le R_{HB}} \frac{\exp\!\left(-\tfrac{1}{2}(\mathbf{x}-\mathbf{m})^{\top} C_2^{-1} (\mathbf{x}-\mathbf{m})\right)}{2\pi\sqrt{|C_2|}} \, \mathrm{d}\mathbf{x} \tag{8.12}$$

The integral is a polar quadrature with 12 Gauss-Legendre points in radius on one interval and 24 midpoint angles, refined to three intervals and 48 angles when the smallest standard deviation of C_2 is below R_{HB}. For a zero miss and an isotropic σ it reproduces $1 - \exp(-R_{HB}^2/2\sigma^2)$ to 10^{-9} for σ from 5 to 1000 m, and for an offset, correlated case it agrees with a 1500×1500 grid within 10^{-4}, two of the application's 43 passing checks. The model stops being valid for slow, co-orbital encounters that last minutes, where the relative motion curves.

> **+ DID YOU KNOW? Ten Kilometers to Be Sure**
>
> In a NASA comparison of conjunctions, the passes with a collision probability above 1 in 10,000 had predicted misses from under 100 m to over 10 km. A rule on the miss alone would need a 10 km limit to catch 95 % of them, and it would also flag more than half of the harmless passes [18].

8.6 Implemented Planners and Lower Layer

No feedback loop closes on the altitude, and this application does not instantiate the framework's velocity-form PID controller of Section 1.6.2. An altitude loop would fire twice an orbit against the J_2 oscillation, and a miss distance is not a quantity a loop can regulate. Every burn is instead decided by one of two planners and executed by a lower layer, all on the operator's side.

8.6.1 Timing

The filter and the lower layer run every 10 s, one physics step (Section 1.6.1), and the order within a step is the controller, the integration, then the post-step updates of the battery, the true mean elements and the true closest approaches. The conjunction planner runs at the screening times above, and the maintenance planner at its decision interval T_{plan}, 6 h by default. One sample per simulated minute goes to the historian, and the app runs from 0.05 to 24 simulated hours per second, one by default.

8.6.2 Maintenance Planner

At each decision the planner reads the filter's mean $\hat{\bar{h}}$ against the band $\bar{h}_{SP} \pm \Delta h$. Above the band it lowers the orbit to $\bar{h}_{SP}$. The deadband mode reboosts to $\bar{h}_{SP} + 0.8\Delta h$ only when $\hat{\bar{h}}$ is already below the band. The predictive mode integrates Equation (8.7) over the horizon H_f in 15 min steps, with $k = \hat{k}$ when the drag scale is estimated. It reboosts to the same target when the forecast crosses $\bar{h}_{SP} - 0.9\Delta h$ within the lead of Equation (8.13).

$$t_{lead} = 2T + \frac{\Delta v}{a_{max}} + t_{slew} + T_{plan}, \qquad \Delta v = \frac{v\,\Delta a}{2a} \tag{8.13}$$

where T is the period of the estimated mean orbit and $a_{max} = \hat{\eta} F / \hat{m}$ when burns are sized with the estimated efficiency. A reboost starts one slew time and one minute after the decision, and with two burns selected each half fires half an orbit apart. The planner then holds for 1.1 orbits after the last burn so the mean settles, and it stays silent while an avoidance or its return is pending.

8.6.3 Avoidance Planner

The planner acts on one pass at a time, the first in the catalog that needs action. The nominal mode acts when the predicted miss is under the floor d_{min} and sizes the burn on the miss alone. The uncertainty-aware mode acts when $P_c > P_{c,max}$ or the miss is under the floor, plans to the target $P^* = P_{c,max}/k_P$, and keeps the floor with $1.5\sigma_{cat}$ in hand, where σ_{cat} is the largest standard deviation of the catalog covariance in the plane. Candidate burns, prograde or retrograde and radial only when allowed, are centered every sixteenth of an orbit from 20 min before TCA back to the end of the slew or to $t_{lead,max}$ before TCA, whichever is later. A candidate is kept only if it starts after the slew and ends at least 20 min before TCA. For each one the linear response of Equation (8.8), averaged over the finite burn by Simpson's rule and resolved on $(\hat{\xi}, \hat{\zeta})$, gives the in-plane displacement $\mathbf{g}$ per unit Δv. With a small disc the probability is $P_c \approx K \exp(-d^2/2)$, where $K = R_{HB}^2/(2\sqrt{|C_2|})$ and d is the Mahalanobis distance of the miss. Both constraints are then quadratic in Δv, and the least velocity change follows in closed form, as shown in Equation (8.14).

$$\Delta v = \frac{b + \sqrt{b^2 - ac}}{a}, \qquad \begin{aligned} a &= \mathbf{g}^\top W \mathbf{g}, \quad b = \mathbf{g}^\top W \mathbf{m} \\ c &= \mathbf{m}^\top W \mathbf{m} - D^2 \end{aligned} \tag{8.14}$$

with $W = I$ and $D = d_{min} + 1.5\sigma_{cat}$ ($D = d_{min}$ in the nominal mode) for the floor, and $W = C_2^{-1}$ and $D^2 = 2\ln(K/P^*)$ for the probability, the larger of the two setting the burn. The cheapest candidate under Δv_{max} is chosen, and with re-screening on it must also leave every other screened object clear on the same linear model. The uncertainty-aware mode then refines its size by bisection on the exact P_c with the burn's own execution error included. Both modes verify the plan by propagating the full model with the finite burn, and a plan that falls short is scaled once, by 1.25 or in the nominal mode by the ratio of floor to miss, at most doubling it. The resulting cost against lead time is shown in Figure 8.2.

When nothing meets the criterion, the planner searches the longest full-thrust burns the clock still allows, up to 2 min before TCA, in both directions, and flies the best of them only if it at least halves the probability (or raises the miss by half in the nominal mode). The pass is then flagged infeasible with its residual probability, whether or not a burn flies.

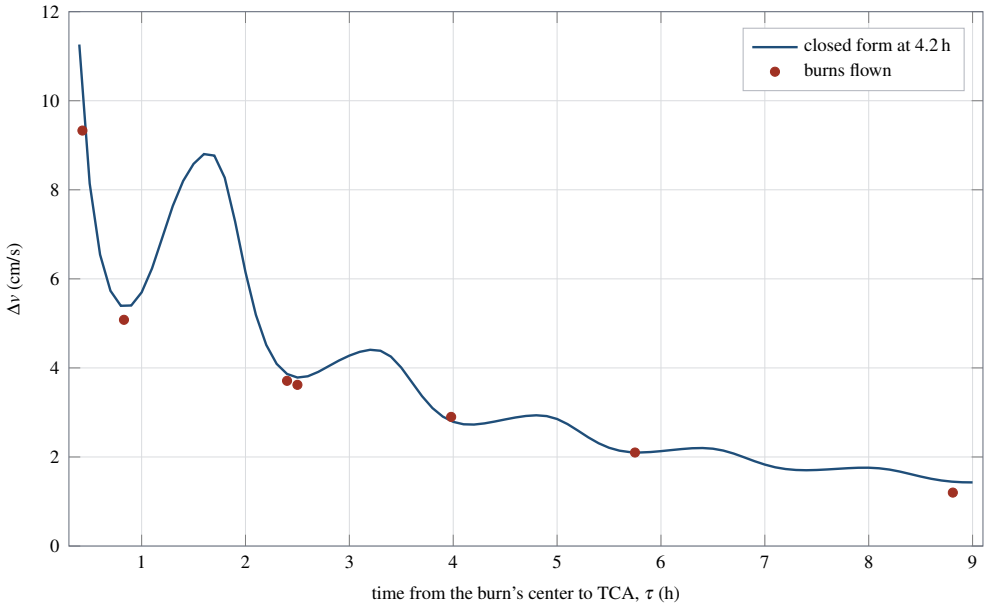


Figure 8.2 Least along-track velocity change that moves Fragment 1 of case 5 to the planner's floor of 272.7 m, against the time from the burn's center to TCA. The curve is Equation (8.14) on the case-5 prediction at 4.2 h, retrograde and averaged over the finite burn. The points are the burns flown with the earliest burn set from 0.5 to 12 h, each planned on the tracking available when it was committed.

8.6.4 Return, Re-Screening and Verification

A return burn, equal and opposite and a whole number of orbits after the avoidance, at least 15 min after TCA, restores the orbit. It is omitted when the avoidance was a prograde burn from below $\bar{h}_{SP}$ that leaves the mean under $\bar{h}_{SP} + 0.8\Delta h$, which is then a free reboost. With re-screening on, the plan and up to three placements of its return are checked against every screened object. Every return or reboost is also screened again on the full model when it is committed, and it is moved one orbit later, at most three times, if it would create a pass over the planning target or under the floor. With verification on, the uncertainty-aware mode re-plans 15 min after an avoidance on the orbit the filter now sees and flies a trim if the pass still needs one. The nominal mode fires and forgets. In every mode the return is resized, until it is committed, from what the filter believes the avoidance and any trim delivered.

8.6.5 Lower Layer and Manual Operation

A burn is committed one slew time before it starts, and its errors are drawn then, $g_b = 1 + \sigma_F \varepsilon$ and a pointing error of σ_θ on each of the two axes across the thrust. It fires at the rated thrust, the last step at the share that completes the impulse, and it pauses when the battery falls to 5 % of E_{bat} and resumes above 25 %. No burn fires once the xenon allocated to the run is spent. The lower layer counts commanded on-time and cannot see the thrust delivered. In manual mode the operator's along-track slider sets ℓ and the sign of the thrust, and two buttons schedule ±5 cm/s burns after the slew. The planners still assess every pass in manual mode, but they plan nothing, and automatic burns wait. A burn whose window passes before the return to automatic expires unflown, and both planners then plan again at once from the orbit as it is. Switching a planner off drops the burns it has planned but not committed.

Table 8.4 Planner and Filter Settings. Controls Are Adjustable, the Rest Fixed in the Source

Setting	Default	Range
maintenance planner	predictive	off, deadband
reboost as	two burns	one or two
horizon H_f, decision T_{plan}	24 h, 6 h	2–72, 1/6–24 h
avoidance planner	aware	off, nominal
$P_{c,max}$	10^{-4}	10^{-6}–10^{-2}
d_{min}, k_P	200 m, 3	0–2000, 1–10
Δv_{max}, $t_{lead,max}$	30 cm/s, 6 h	1–100, 0.5–24
re-screen, verify	on, on	switches
radial burns	off	advanced switch
size with $\hat{\eta}$, estimate $\hat{k}$	on, on	switches
$\hat{\rho}_s$, q_k	1, 0.06 $\mathrm{h}^{-1/2}$	0–5, 0–0.6
screening interval	30 min	fixed
final assessment	TCA − 10 min	fixed
last end of a planned burn	TCA − 20 min	fixed
floor margin, spacing	$1.5\sigma_{cat}$, $T/16$	fixed
battery pause, resume	5, 25 %	fixed
q_a ($\mathrm{m/s}^{1.5}$), q_η ($\mathrm{s}^{-1/2}$)	2×10^{-9}, 10^{-6}	fixed

Table 8.5 Scenario Sequence and the Changes From the Defaults

#	Scenario and change	Teaching point
1	coast, $s = 0$, $\hat{\rho}_s = 0$	mean against altitude
2	decay, both planners off	120 m a day
3	$\bar{h}_0 = 499.80$ km, deadband every 10 min, one burn	the far side rises
4	$\bar{h}_0 = 499.95$ km, deadband, storm at 6 h	a storm between decisions
5	Fragment 1, late fragment	lead time buys the pass
6	rocket body, nominal planner	distance is not risk
7	Fragment 2, $\eta = 55$ %, verification off	only the orbit shows
8	Fragment 3, Satellite B, re-screening off	the return makes a pass
9	96 h campaign, 8 g allocated	graded challenge

8.6.6 Truth, Operator and Analysis Models

The truth model is Equations (8.1)–(8.5) with the true density scale, the storm, the true thruster output and the true debris. The operator's model is the same force model with $\hat{k}\hat{\rho}_s$ for the density, the commanded thrust times $\hat{\eta}$ along the nominal direction, and the catalog, plus the linear relative motion for sizing burns and carrying covariances and Equation (8.7) for the maintenance forecast. The analysis model is Equations (8.6), (8.8) and (8.9), which this chapter uses to interpret the runs. The settings of the planners and the filter are listed in Table 8.4.

8.7 Cases and Measured Outcomes

Table 8.5 lists the eight cases and the campaign, and Table 8.6 their objects. Each case starts from the defaults of Table 8.4 except for the changes in Table 8.5. The numbers below come from headless engine runs at the default seed, in hours of simulated time, and a configuration gives one run whether it is set before the run starts or carried through Reset (Section 1.6.4). A pass counts as resolved when its final assessment gives $P_c \leq 10^{-4}$ and a predicted miss of at least 200 m. These thresholds, like the 250 m band, are the app's design choices, and Table 8.7 collects the outcomes.

Table 8.6 Synthetic Objects. Miss Vector (ξ, ζ) in Meters, Standard Deviations in the Object's Radial, Along-Track and Cross-Track Axes

Case, object	Alert, TCA (h)	θ (°)	Miss	σ_0 (m)	σ_F (m)	N_u
5, Fragment 1	0.5, 10.2	160	60, −40	40/300/40	15/80/15	4
5, late fragment	20, 20.42	150	−15, 8	80/400/80	60/250/60	1[a]
6, rocket body	0.5, 14.1	100	120, 700	300/3000/300	150/1500/150[b]	4
7, Fragment 2	0.5, 10.2	160	60, −40	40/300/40	15/80/15	4
8, Fragment 3	0.5, 10.2	160	60, −40	40/300/40	15/80/15	3[c]
8, Satellite B	4.86, 15.39[d]	90	3, 2	6/30/6	4/20/4[e]	3
9, Fragment 1	0.5, 10.0	160	60, −40	40/300/40	15/80/15	4
9, rocket body	38, 52.0	100	120, 700	300/3000/300	150/1500/150[b]	4
9, Fragment 2	62, 72.0	150	−50, 30	40/300/40	15/80/15	4

[a] last update 0.25 h before TCA. [b] true offset drawn from 20/100/20. [c] last update 6.2 h before TCA. [d] catalogued 0.5 h after Fragment 3's avoidance and placed three orbits after its planned return. [e] true offset drawn from 2/5/2.

Table 8.7 Results of the Cases in Headless Runs at Seed 20261102

Case	Measured outcome
1	$\lvert\mathbf{r}\rvert - R_E$ 498.42–501.58 km, $\bar{h}$ within 0.98 m for 24 h
2	$\bar{h}$ 499.885, 499.763, 499.643 km at 24, 48, 72 h, out at 50.79 h
3, one burn	24.91 cm/s at 10.94 h, perigee 499.770, apogee 500.575 km
3, two burns	2×12.45 cm/s, perigee 500.176, apogee 500.179 km
4, deadband	lowest 499.677 km, 6.08 h outside, 2.52 g
4, predictive	reboost 12.10 h, lowest 499.848 km, none outside, 1.68 g
5, Fragment 1	2.10 cm/s at 4.43 h, final 245 m, true 278 m
5, late, best effort	18.11 cm/s, P_c $1.6 \times 10^{-2} \to 1.1 \times 10^{-3}$, infeasible
5, nominal	1.46 cm/s, final 171 m (true 204), unresolved
6, nominal	no burn, final 1622 m, $P_c = 1.8 \times 10^{-4}$, true 807 m
6, aware	3.66 cm/s at 8.12 h, final 3209 m, $P_c = 2.5 \times 10^{-6}$
7, no verification	final 138 m (true 171), unresolved, 4.20 cm/s
7, verification	trim 1.18 cm/s, final 232 m (true 265), 7.96 cm/s
8, unscreened	B at 3.1 m, $P_c = 0.23$, then 0.62 cm/s for B, 4.35 cm/s
8, re-screened	return one orbit later, B at 216 m, 3.73 cm/s

8.7.1 The Mean Orbit and the Decay

The coast of case 1 and the decay of case 2 show the plant before any planner acts. The first shows the 1.6 km oscillation of the altitude against a mean that holds to 0.98 m, and the second a mean that falls to 499.885, 499.763 and 499.643 km at 24, 48 and 72 h and leaves its band at 50.79 h. The decay is nearly steady because 360 m is a small part of the 64 km scale height at 500 km. A controller chasing the altitude of case 1 would fire twice an orbit to cancel an oscillation that costs nothing.

8.7.2 A Reboost and a Storm

Case 3 starts 50 m above the bottom of the band with a deadband deciding every ten minutes. The true mean crosses 499.75 km at 10.84 h, and the decision at the start of that 10 s step already finds the filter's mean 2 cm below the edge, so 24.91 cm/s fires from 10.94 h for 24.9 min, 2.18 g of xenon. At 16 h the mean stands at 500.173 km, but the perigee stays at 499.770 km, where the burn fired, and the apogee has risen to 500.575 km, 325 m above the band. Split into two burns of 12.45 cm/s at 10.94 and 11.72 h, the reboost leaves the perigee at 500.176 km and the apogee at 500.179 km, an eccentricity of 2.6×10^{-7}. The mean spends 0.36 h below the band either way, because the burn starts a slew time and a minute after the decision and the mean lags half an orbit behind it. The filter's mean is so close to the edge at that decision that on five of six other seeds the reboost waits for the 11.00 h decision and the mean spends 0.53 to 0.54 h below the band.

> **TRY IT IN THE APP 3 · Scheduled Reboost**
>
> Load the case and run to 16 h. On the Mean-orbit perigee and apogee trend the single reboost leaves the perigee at 499.770 km, where the burn fired, and lifts the apogee to 500.575 km. Press Reset, set Reboost as to Two burns before pressing Run, and run again: the perigee and apogee end at 500.176 and 500.179 km.

Case 4 triples the density from 6 h, and the filter's drag scale passes 2.5 at 8.72 h from the navigation fixes alone. With maintenance off the decay steepens to 361 m a day between 10 and 20 h. The deadband decides every 6 h. At its 18 h decision the mean reads 499.761 km, 11 m above the edge, so it waits. The mean crosses at 18.76 h, and the reboost, two burns of 14.42 cm/s, comes only at the 24 h decision. The mean bottoms at 499.677 km, 73 m below the band, and spends 6.08 h outside. The predictive planner forecasts a crossing at 17.25 h at its 12 h decision, reboosts at 12.10 h with two burns of 9.63 cm/s, keeps the mean above 499.848 km and uses 1.68 g against 2.52 g. The deadband's failure comes from waiting, not from its target.

8.7.3 Lead Time and a Late Alert

Case 5 warns at 0.5 h of Fragment 1, predicted 70.7 m away at 10.20 h, nearly head-on at 15.0 km/s, with $P_c = 1.0 \times 10^{-2}$. The uncertainty-aware planner, allowed to burn no earlier than 6 h before TCA, flies 2.10 cm/s retrograde from 4.43 h for 126 s, and a matching return after the pass. The final prediction is 245 m with a probability below 10^{-12}, and the truth revealed after the pass is 278 m. The state of this case at 4.2 h, before the burn is committed, is shown in Figure 8.3. Allowed 12 h, the planner flies 1.20 cm/s centered 8.81 h ahead. Held to 3 h it needs 3.62 cm/s, held anywhere from 2 to 0.9 h it needs 5.08 cm/s centered 0.83 h ahead, about half an orbit, and held to 0.5 h it needs 9.33 cm/s, as shown in Figure 8.2. The closed-form curve falls roughly as $1/\tau$ with a ripple of one orbital period, its minima about one orbit apart and the first near half an orbit, where the radial displacement of Equation (8.8) is largest. The points depart from the curve because each burn was planned on the tracking of its own moment.

Pause. Try Exercise 8.1 before reading on.

At 20 h a late fragment is reported 25 min before a pass predicted at 17.1 m. With the 5 min slew and no planned burn allowed in the last 20 min, nothing meets the criterion. The best effort is 18 min of full thrust, 18.11 cm/s prograde, ending 2 min before TCA, and it cuts the final probability from the 1.6×10^{-2} of a pass with no burn to 1.1×10^{-3}, eleven times the threshold. The pass is flagged infeasible under both planners. The nominal planner handles Fragment 1 with 1.46 cm/s sized to the 200 m floor on the prediction of its moment, and ends at a predicted 171 m, inside the floor, once later tracking moves the fragment.

8.7.4 A Miss Distance That Hides Its Risk

Case 6 predicts a rocket body 725 m away at 14.1 h with a 3 km along-track standard deviation that tracking brings to 1.5 km. The crossing angle is 100°, so Equation (8.9) puts 0.643 of that error into the plane. At the final assessment the combined ellipse is 971 by 150 m, the predicted miss 1622 m, eight times the floor, and $P_c = 1.8 \times 10^{-4}$, over the threshold. The nominal planner flies nothing. The rocket body passes 807 m away, 839 m from where it was predicted. The uncertainty-aware planner burns 3.66 cm/s prograde at 8.12 h, and the pass ends at a predicted 3209 m with $P_c = 2.5 \times 10^{-6}$ and a true miss of 2386 m. The miss stays within 2° of the long axis of the ellipse in both runs, and the burn doubles its Mahalanobis distance from 1.71 to 3.38.

Pause. Try Exercise 8.2 before reading on.

The same case shows dilution in numbers. At the alert the ellipse is 2010 by 301 m and $P_c = 1.6 \times 10^{-4}$, while scaling that covariance would give at most 4.7×10^{-4}, with its standard deviations at 0.38 of their size. As tracking shrinks the ellipse to 1155 by 178 m the probability rises to 2.8×10^{-4}, and it falls to 1.8×10^{-4} only when the last update moves the predicted miss out to 1622 m. At that

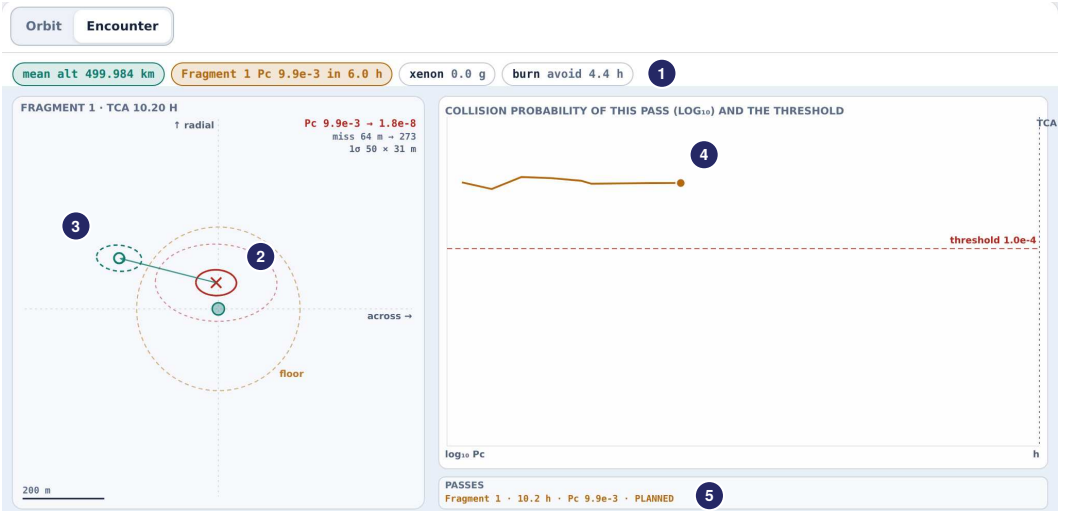


Figure 8.3 The application in case 5, a well-observed conjunction, at $t = 4.20$ h with the shipped settings, 6.0 h before Fragment 1 reaches its closest approach at 10.20 h, cropped to the plant view. (1) The badges give the next burn as an avoidance at 4.4 h. (2) The Encounter view shows the plane of the pass, centered on the satellite and its 15 m hard-body circle with radial up, with the fragment 64 m away inside its combined 1σ ellipse of 50 by 31 m and its 3σ ellipse, and the dashed 200 m floor. (3) An arrow leads to the ring where the planned burn will put it, 273 m away inside its own dashed ellipse. The readout gives P_c from 9.9×10^{-3} now to 1.8×10^{-8} with the plan. (4) Beside the plane, the probability of this pass since the alert at 0.5 h stays between 6.5×10^{-3} and 1.5×10^{-2}, two decades above the 10^{-4} threshold, and (5) the list of passes marks Fragment 1 as planned.

point the probability is 92 % of the largest any scaling of the covariance gives, reached at 1.21 times its standard deviations. A covariance larger than the worst case makes a pass look safer than better data will, the behavior analyzed in [7].

8.7.5 A Thruster That Underdelivers

Case 7 repeats Fragment 1 with the thruster delivering 55 % of its command and verification off. The 2.10 cm/s burn delivers 0.548 of it. The filter's efficiency reads 0.82, 0.68 and 0.63 at 8, 14 and 22 min after the burn ends. Fired and forgotten, the pass ends at a predicted 138 m, inside the floor, with a true miss of 171 m. With verification on, the planner re-screens 15 min after the burn on the orbit the filter now sees and flies a trim at 5.02 h sized to deliver 1.18 cm/s with the estimated efficiency. The pass ends at 232 m (true 265 m), and the return is sized from what the filter believes both burns delivered, 7.96 cm/s commanded in all against 4.20 cm/s. The xenon flow of Equation (8.3) follows the command, so the weak thruster spends the full propellant for every burn.

8.7.6 A Return Burn That Creates a Pass

Case 8 avoids Fragment 3 with 1.86 cm/s at 4.33 h and plans an equal return at 10.64 h. Half an hour after the avoidance the catalog lists Satellite B, an active spacecraft with a shared ephemeris, predicted 670 m away at 15.39 h. The world placed it three orbits after the planned return, so it is invisible on the operational trajectory, which holds only committed burns. Committed unscreened at 10.56 h, the return brings B to a predicted 8.4 m, and to 3.1 m with $P_c = 0.23$ once it has flown. With the planner switched off at that moment the true miss would be 14.8 m, inside the hard-body radius, and the run logs a collision. With the planner on and 4.8 h to go, a third burn of 0.62 cm/s clears B

Table 8.8 Campaign Configurations Over 96 h. Out Is the Time the True Mean Spent Outside the Band

Configuration	Unres.	Worst P_c	Out (h)	Xenon (g)	Burns
shipped, deadband, nominal	3	1.7×10^{-3}	5.53	2.83	7
predictive, nominal[b]	3	2.1×10^{-3}	0	4.16	10
deadband, aware	0	4.9×10^{-5}	5.51	4.59	9
deadband, planner off	3	2.4×10^{-2}	5.59	2.35	1
suggested[a]	0	7.7×10^{-5}	0	6.04	12
predictive one burn, aware	0	4.9×10^{-5}	0	5.96	10

[a] predictive, two burns, uncertainty-aware, 69.1 cm/s in all

[b] two-burn reboosts

to 241 m, 4.35 cm/s in all. With re-screening on, the return is screened before it is committed and moved one orbit later, and B passes 216 m away with no extra burn, 3.73 cm/s in all.

8.7.7 The Four-Day Campaign

The campaign decays the orbit from 500.05 km, triples the drag from 14 h, sends Fragment 1 at 10 h, the rocket body at 52 h and, after the thruster falls to 55 % at 58 h, Fragment 2 at 72 h, with 8 g of xenon allocated. Its target is every pass resolved, no time with the true mean outside the band, the mean in the band at the end and the xenon within the allocation. The six configurations measured appear in Table 8.8, and Figure 8.4 draws three of its columns as bars. In its delivered form, a one-burn deadband with nominal avoidance leaves all three passes unresolved, Fragment 2 at a predicted 76 m, with a worst final P_c of 1.7×10^{-3}, 17 times the threshold. It spends 5.53 h outside the band and 2.83 g of xenon. In the configurations measured, predictive maintenance is what keeps the band and the uncertainty-aware planner is what resolves the passes, so the two failures have separate remedies. With both, the worst final probability is 7.7×10^{-5}, 23 % under the threshold, for the rocket body at 1885 m, and Fragment 2 needs two trims after the thruster fault. The smallest final miss is 212 m, 12 m over the floor, the lowest mean is 499.800 km, 50 m inside the band, and the four days use 6.04 g of the 8 g. A one-burn predictive reboost also meets the target, because the target counts the mean and not the apogee. Over six other seeds, nominal avoidance leaves one to three passes unresolved and the deadband 5.47 to 9.00 h outside the band, while the uncertainty-aware planner resolves every pass and the suggested configuration meets the target each time with 3.22 to 5.23 g of xenon.

> ✦ DID YOU KNOW? **Safety Costs Xenon**
>
> From the simulator: in the four-day campaign the shipped configuration spends 2.83 g of xenon in 7 burns and leaves all three passes unresolved. The suggested configuration, which meets the target, spends 6.04 g in 12 burns, more than twice as much (Table 8.8).

8.8 Technology in Practice

The satellite of Section 8.3 is read by navigation fixes and a debris catalog, and it is moved by one Hall thruster. Figure 8.5 draws its loop with real devices in the place of the schematic's symbols.

What Measures It

The fixes of Equation (8.10), 5 m of white noise on each axis every 10 s, stand for a satellite navigation receiver, now the main means of orbit determination in low orbit [11]. Nothing aboard measures the mean altitude, the drag or the thrust delivered, and the filter infers all three from the fixes. An operator's own prediction of its orbit is often the best one available, because the operator holds the receiver's fixes, a good model of the drag area and the planned maneuvers. The handbook of practice

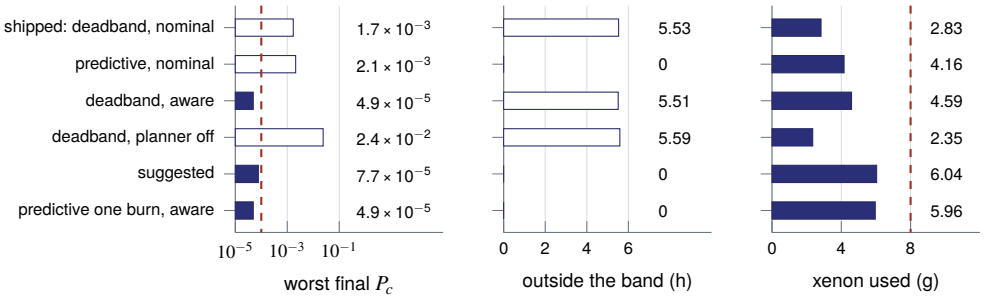


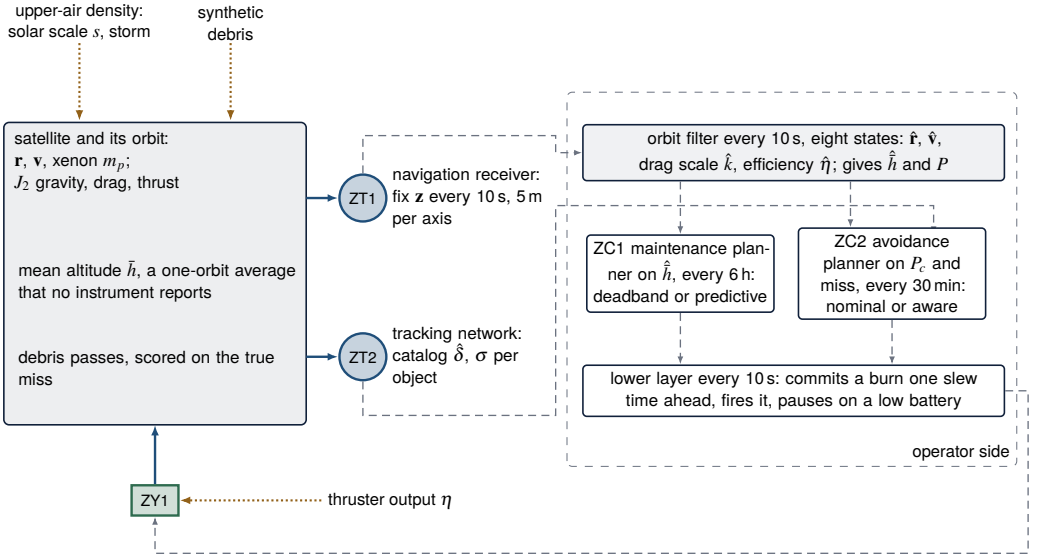
Figure 8.4 Three columns of Table 8.8 drawn as bars for the six configurations of the four-day campaign, with its limits as dashed lines: the 10^{-4} threshold on the worst final collision probability (logarithmic axis) and the 8 g of xenon allocated. Solid bars meet the target of their panel, which for the time outside the band is none, and white bars miss it. Notice that the passes are resolved only by the uncertainty-aware planner and the band is kept only by predictive maintenance, so only the configurations with both meet the target, and they spend the most xenon.

asks that every point of that prediction carry a realistic covariance [18], the role of the satellite's own covariance at TCA in Equation (8.11). The catalog of the debris, with its declared standard deviations, stands for the tracking network that the second history segment describes.

> **TECHNOLOGY** **Sensor: Satellite Navigation Receiver**
>
> For a satellite in low orbit, a receiver of the Global Positioning System (GPS) is now the main means of orbit determination, in place of tracking from the ground. It times the signals of at least four navigation satellites: three fix the position and the fourth corrects the receiver's clock. Its accuracy is limited by the passage of the signals through the upper atmosphere and by the precision of the civil code. Jamming and spoofing of these signals have been detected in orbit, and a satellite should be able to tolerate their loss [11]. On a trend a lost receiver is a gap in the fixes, across which the filter runs on its model alone.

What Moves It

The thruster of Table 8.2, 20 mN at 350 W with a specific impulse of 1400 s, gives little thrust for its power, so its burns are long. A reboost of 24.91 cm/s fires for 24.9 min, more than a quarter of an orbit, which is why the avoidance planner averages each burn's response over its length. The magnitude and pointing errors drawn for each burn in the lower layer stand for the performance drift and the walk of the thrust vector of a real thruster [11], and the output η of case 7 for a thruster that has lost part of its thrust. The model has no attitude, so the slight roll torque that a Hall thruster puts on the spacecraft over a long firing, which reaction wheels may have to absorb [11], is left to Chapter 7.



ZY1: Hall thruster drive, 20 mN at 1400 s on xenon, 350 W; the lower layer counts commanded on-time and cannot see the thrust delivered; each burn carries a 2 % magnitude and a 1° pointing error.

Figure 8.5 The loop of the chapter with real devices in place of the schematic's symbols. ZT1 is the satellite navigation receiver and ZT2 the tracking network behind the debris catalog. On the operator's side (dashed frame) the orbit filter feeds the maintenance planner (ZC1) and the avoidance planner (ZC2), which also reads the catalog, and the lower layer flies their burns on the Hall thruster drive ZY1. Notice that neither controlled quantity is measured: the mean altitude is an average of the filter's estimate and the collision probability a prediction for each pass.

TECHNOLOGY Final Element: Hall Thruster

A Hall thruster sets a strong axial electric field and a radial magnetic field at the exit of its discharge chamber. Electrons held in a current around the chamber ionize the incoming propellant, the ions are accelerated by the electric field, and a cathode neutralizes the plume. Electric thrusters give less than about 75 mN of thrust for each kilowatt, so they may fire for hundreds of hours where a chemical thruster needs seconds or minutes, which suits station keeping. Erosion of the plasma-wetted surfaces changes the performance over the thruster's life, and the thrust vector can walk while the thruster warms up [11]. A weak thruster shows only in the orbit, as in case 7.

Where the Controller Runs

The chapter's filter and lower layer run every 10 s, the avoidance planner every 30 min, and the maintenance planner every 6 h. In operations an operator's flight dynamics system produces the predicted orbit with its planned maneuvers and sends it at least daily to a screening authority, which checks it against the catalog [11]. The structure of this chapter therefore lives in decision layers above any loop, not in a function block of a controller (Section 8.6).

What Is Modernizing It

Collision avoidance is being tried on board. A NASA swarm demonstration launched in 2023 carried software that computes the probability of collision and proposes maneuvers to reduce the risk, and a ground hub is planned to screen such proposed trajectories in near real time and return the results to every operator involved [11]. Hall thrusters fed with krypton, less efficient than xenon but much cheaper, are expected to become more common, above all for large constellations [11]. The

planner itself is a chance-constrained decision on a carried-forward uncertainty, the form of stochastic predictive control [8], and the uncertainty-aware planner of Section 8.6 is a small instance of it.

> **IN PRACTICE** **A Threshold, a Margin and a Deadline**
>
> NASA's best practices for conjunction assessment record a broad convergence on mitigating a pass when its probability of collision exceeds 1 in 10,000, and also when the predicted miss is smaller than the hard-body radius. A mitigation should lower the probability by at least 1.5 orders of magnitude and create no new conjunction above the threshold, which is checked by screening the planned trajectory again. The decision falls due at a maneuver commitment point before the closest approach. An earlier burn can be smaller, while waiting may show that no burn is needed [18].

8.9 Limitations

The debris is synthetic, placed to produce each pass, with covariances declared rather than derived from tracking, and the atmosphere is a table times two scales. No number in the app is an orbit forecast or a statement about any real object, and its thresholds are design choices rather than validated practice.

The collision probability is the short-encounter model with independent Gaussian errors, and the satellite's covariance is carried with the linear relative motion, which ignores the effect of J_2 on that motion. Burns are sized on the same linear model and verified on the full one, which corrects their size but not the covariance.

The satellite has no attitude dynamics, navigation is position only, and the loop scorecard of the framework, computed for the mean altitude, means little for a plant flown by planners. This chapter therefore reports passes, times, altitudes and xenon instead. A burn already committed when its planner is switched off still flies, because its slew has begun, and the challenge score reads the mean altitude only, so it cannot tell a round orbit from an eccentric one.

> **OPEN CHALLENGE** **Covariances Worth Trusting**
>
> A probability of collision is only as good as the covariances that feed it. NASA's handbook found the realism of the covariances in operators' own predicted orbits much more uneven than in the catalog, with few published studies and no practical guides for testing it. Drag adds an error that is hard to bound, because the space weather indices that drive density models are difficult to predict [18]. In case 6 the probability rose as tracking shrank the ellipse. Progress would be covariances whose realism is tested routinely against later tracking, with the density forecast error carried inside them.

8.10 Conclusion

The simulator of this chapter flies a low Earth orbit satellite on a running mean of its altitude and the collision probability of each pass, both kept apart from the truth that produces them. With gravity alone the altitude swings 1.6 km twice an orbit while the mean holds to 0.98 m, and drag lowers that mean about 120 m a day. One reboost burn lifts the apogee 325 m above the band, and a deadband deciding every six hours spends 6.08 h below it through a storm that a predictive planner never lets the mean reach. A 71 m pass costs 1.20 cm/s with 8.8 h of lead, 2.10 cm/s with 5.8 h and 5.08 cm/s with less than an orbit, while 18.11 cm/s with 25 min of warning leaves a probability eleven times the threshold. A rocket body predicted 1.6 km away carries $P_c = 1.8 \times 10^{-4}$ because its ellipse is 971 by 150 m. A 55 % thruster leaves a pass at 138 m unless the planner checks the orbit and trims, and an unscreened return brings an active satellite to 3.1 m. The model stops at J_2 with a tabulated atmosphere, synthetic debris and a linear own covariance. A fuller model would add a density forecast with its errors, catalog-wide screening, attitude and slew dynamics and velocity measurements in the filter. The application is a browser plant on which a satellite is flown, as in operations, on what its operators predict rather than on what they measure.

KEY POINTS

- A satellite in low orbit is flown on the average of its height over one orbit and on the odds of each collision, and no instrument reports either.
- With gravity alone the altitude swings 1.6 km twice an orbit while the mean holds to 0.98 m, and drag lowers that mean about 120 m a day.
- A deadband deciding every six hours spends 6.08 h below the band through a storm that a predictive planner never lets the mean reach.
- A rocket body predicted 1.6 km away carries $P_c = 1.8 \times 10^{-4}$, over the threshold, because its ellipse is 971 by 150 m.

Where else this applies. Volume 9, Chapter 7 commits its releases days ahead on an uncertain forecast under a chance constraint, and Volume 5, Chapter 13 steers on a filter's estimate of a position nothing aboard measures.

References

[1] Y. Kozai, "The motion of a close earth satellite," *The Astronomical Journal*, vol. 64, p. 367, 1959.

[2] D. G. King-Hele and D. M. C. Walker, "Density of the upper atmosphere determined from satellite orbits, and its variation during 1957–61," *Nature*, vol. 191, no. 4784, pp. 114–118, 1961.

[3] J. T. Emmert, "Thermospheric mass density: A review," *Advances in Space Research*, vol. 56, no. 5, pp. 773–824, 2015.

[4] M. R. Akella and K. T. Alfriend, "Probability of collision between space objects," *Journal of Guidance, Control, and Dynamics*, vol. 23, no. 5, pp. 769–772, 2000.

[5] F. K. Chan, *Spacecraft Collision Probability*. American Institute of Aeronautics and Astronautics, Inc., 2008.

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

[7] M. S. Balch, R. Martin, and S. Ferson, "Satellite conjunction analysis and the false confidence theorem," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 475, no. 2227, p. 20180565, 2019.

[8] M. Farina, L. Giulioni, and R. Scattolini, "Stochastic linear Model Predictive Control with chance constraints – a review," *Journal of Process Control*, vol. 44, pp. 53–67, 2016.

[9] T. Flohrer, H. Krag, and H. Klinkrad, "ESA's process for the identification and assessment of high-risk conjunction events," *Advances in Space Research*, vol. 44, no. 3, pp. 355–363, 2009.

[10] O. Montenbruck and E. Gill, *Satellite Orbits*. Springer Berlin Heidelberg, 2000.

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

[12] T.-W. Fang, A. Kubaryk, D. Goldstein, Z. Li, T. Fuller-Rowell, G. Millward, H. J. Singer, R. Steenburgh, S. Westerman, and E. Babcock, "Space weather environment during the SpaceX Starlink satellite loss in February 2022," *Space Weather*, vol. 20, no. 11, p. e2022SW003193, 2022.

[13] A. Garulli, A. Giannitrapani, M. Leomanni, and F. Scortecci, "Autonomous low-earth-orbit station-keeping with electric propulsion," *Journal of Guidance, Control, and Dynamics*, vol. 34, no. 6, pp. 1683–1693, 2011.

[14] D. J. Kessler and B. G. Cour-Palais, "Collision frequency of artificial satellites: The creation of a debris belt," *Journal of Geophysical Research: Space Physics*, vol. 83, no. A6, pp. 2637–2646, 1978.

[15] T. Wang, "Analysis of debris from the collision of the Cosmos 2251 and the Iridium 33 satellites," *Science & Global Security*, vol. 18, no. 2, pp. 87–118, 2010.

[16] C. D. Bussy-Virat, A. J. Ridley, and J. W. Getchius, "Effects of uncertainties in the atmospheric density on the probability of collision between space objects," *Space Weather*, vol. 16, no. 5, pp. 519–537, 2018.

[17] C. Bombardelli and J. Hernando-Ayuso, "Optimal impulsive collision avoidance in low earth orbit," *Journal of Guidance, Control, and Dynamics*, vol. 38, no. 2, pp. 217–225, 2015.

[18] National Aeronautics and Space Administration, "NASA spacecraft conjunction assessment and collision avoidance best practices handbook," National Aeronautics and Space Administration, Washington, DC, Special Publication NASA/SP-20205011318, Dec. 2020. [Online]. Available: https://ntrs.nasa.gov/citations/20205011318

8.11 Exercises

8.1 Concept check · 2 min

In case 5 the uncertainty-aware planner clears Fragment 1 with 1.20 cm/s when its burn may be centered 8.8 h before the closest approach, and needs 5.08 cm/s when it is held to under one orbit. What chiefly explains the difference?

(a) An along-track burn changes the period, so the shift it buys along the track grows with every second before the pass

(b) The catalog covariance is smaller early in the alert, so an early burn has a smaller miss to reach

(c) The thruster delivers more of its command early in the run, while the battery is still full

(d) Drag acts on the lowered orbit for longer and multiplies the effect of an early retrograde burn

8.2 Concept check · 2 min

At the final assessment of case 6 the rocket body is predicted to miss by 1622 m, eight times the 200 m floor. The combined 1σ ellipse in the encounter plane is 971 by 150 m. Which statement is right?

(a) The pass is safe, because the predicted miss is larger than the whole 1σ ellipse

(b) The pass needs a burn, because the miss lies near the long axis, only 1.71 standard deviations out

(c) The pass is safe, because more tracking would shrink the ellipse and could only lower the probability

(d) The pass needs a burn, because the hard-body radius is close to the short axis of the ellipse

8.3 Concept check · 1 min

In case 3 a single prograde reboost of 24.91 cm/s fires while the mean altitude is at the bottom of its band. Once the mean has settled, where is the orbit higher than before?

(a) Everywhere by the same amount, because the mean has risen by about 450 m

(b) At the point where the burn fired, which rises by about twice the change in the mean

(c) Half an orbit from the burn, which rises by about twice the change in the mean

(d) Nowhere until a second burn fires, because one burn changes only the eccentricity

8.4 Theory and calculation · 6 min

The satellite of this chapter has a mass of 120 kg, a Hall thruster of 20 mN at a specific impulse of 1400 s and a mean orbit of radius $a = R_E + 500\,\mathrm{km} = 6878.137\,\mathrm{km}$, with $\mu = 3.986004418 \times 10^{14}\,\mathrm{m^3/s^2}$ and $g_0 = 9.80665\,\mathrm{m/s^2}$.

(a) Find the change in the mean radius per m/s of along-track velocity change, and the rise of the mean from the 24.91 cm/s reboost of case 3. *Checkpoint:* 1.807 km per m/s and 450 m.

(b) Find the time this burn takes at the rated thrust and the xenon it uses. *Checkpoint:* 1495 s (24.9 min) and 2.18 g.

(c) An impulse along the track makes the perigee-to-apogee difference $2\Delta a$. A burn spread evenly over an arc φ of the orbit, centered on the same point, makes it smaller by the factor $\sin(\varphi/2)/(\varphi/2)$. With a period of 94.62 min, estimate the difference this burn leaves and compare it with the mean perigee and apogee that case 3 reports. *Checkpoint:* about 801 m, against 805 m measured.

8.5 Theory and calculation · 7 min

At the final assessment of case 6 the combined 1σ ellipse of the rocket body is 971 by 150 m, the predicted miss lies at a Mahalanobis distance $d = 1.71$, and the combined hard-body radius is $R_{HB} = 15\,\mathrm{m}$. Use the small-disc form $P_c \approx K \exp(-d^2/2)$ with $K = R_{HB}^2/(2\sigma_1\sigma_2)$.

(a) Compute the collision probability. *Checkpoint:* 1.79×10^{-4}.

(b) Scale both standard deviations by a factor s, which divides d by s and K by s^2. Find the s that makes the probability largest and that largest value, and the share of it that the actual covariance gives. *Checkpoint:* $s = 1.21$, 1.94×10^{-4} and 92 %.

(c) The uncertainty-aware planner plans to $P^* = P_{c,max}/k_P$ with $P_{c,max} = 10^{-4}$ and $k_P = 3$. Find the Mahalanobis distance that meets P^* on this covariance. *Checkpoint:* 2.51.

8.6 Theory and calculation · 6 min

An along-track velocity change of 1 cm/s is made at 500 km, where the mean motion is $n = 1.1068 \times 10^{-3}$ rad/s. Use Equation (8.8) for the radial displacement x and the along-track displacement y a time τ after the burn.

(a) Find x and the size of y half an orbit after the burn ($n\tau = \pi$) and one orbit after it ($n\tau = 2\pi$). *Checkpoint:* 36.1 m and 85.2 m; 0 m and 170.3 m.

(b) The pass crosses at 160°, so the radial displacement enters the encounter plane whole and the along-track one multiplied by $\cos(\theta/2)$ (Equation (8.9)). Find the in-plane displacement at both times, and say which lead the burn should prefer. *Checkpoint:* 39.0 m at half an orbit, 29.6 m at one orbit.

8.7 App challenge · 8 min

Load scenario 4, Density surge, and change nothing: the maintenance planner is a deadband deciding every 6 h, and a storm triples the density from 6 h. Use the default seed and the default speed, one simulated hour per second.

(a) Run to 48 h. Read the Time outside the band and the Propellant used tiles, and find in the exported file the lowest mean altitude, column `h_mean_km`, and the first row at which the filter's drag scale, column `drag_scale_est`, exceeds 2.5. *Checkpoint:* 6.08 h, 2.5 g; 499.677 km at 24.3 h; 8.73 h, within 0.02 h.

(b) Press Reset, set Maintenance planner to Predictive before pressing Run, and run to 48 h again. Read the same tiles and the lowest mean altitude. *Checkpoint:* 0.00 h, 1.7 g; 499.848 km.

Accept altitudes within 0.002 km. Explain which part of the deadband's rule lets the mean leave the band.

8.8 App challenge · 8 min

Load scenario 5, Well-observed conjunction. Use the default seed and the default speed. For each run read the Δv commanded tile at any time between 9.5 and 10.6 h, after the avoidance burn and before the return burn is committed, and in the exported file the predicted miss, column `miss_pred_m`, in the row at 10.10 h, after the final assessment.

(a) Run as shipped, with Earliest burn before TCA at 6 h, and pause at about 10.6 h. *Checkpoint:* 2.1 cm/s and 245 m.

(b) Press Reset, set Earliest burn before TCA to 12 h before pressing Run, and run to 10.6 h. *Checkpoint:* 1.2 cm/s and 204 m.

(c) Press Reset, set Earliest burn before TCA to 1.5 h, and run to 10.6 h again. *Checkpoint:* 5.1 cm/s and 231 m.

Accept a miss within 1 m. Column `dv_cmd_mps` gives the burn in m/s to more digits.

8.9 App challenge · 8 min

Load the Challenge, four days in the band. Set the speed to 12×, so that each 96 h run takes 8 s, use the default seed, and read the Unresolved passes, Worst P_c at TCA, Time outside the band and Propellant used tiles when the run stops at 96 h. Make every change before pressing Run.

(a) Run the challenge as shipped: one-burn deadband maintenance and nominal avoidance. *Checkpoint:* 3 of 3, 1.7×10^{-3}, 5.53 h, 2.8 g.

(b) Press Reset, set Avoidance planner to Uncertainty-aware and run again. *Checkpoint:* 0 of 3, 4.9×10^{-5}, 5.51 h, 4.6 g.

(c) Press Reset, also set Maintenance planner to Predictive and Reboost as to Two burns, and run again. *Checkpoint:* 0 of 3, 7.7×10^{-5}, 0.00 h, 6.0 g.

Say which change cures which failure, and whether the target is met.

8.10 Application area · 4 min

A mission planner sizes the xenon of this satellite for the default atmosphere, in which drag lowers the mean altitude by an orbit average of 120.4 m a day at 500 km. Use $a = 6878.137$ km, $v = 7612.6$ m/s and 87.4 mg of xenon per cm/s.

(a) Find the velocity change per year that holds the mean altitude. *Checkpoint:* about 24.3 m/s a year.

(b) Find the xenon per year and the years that the 2.0 kg tank lasts, and name two effects of the Limitations section that make this an upper bound on the lifetime. *Checkpoint:* about 213 g a year and 9.4 years.

8.11 Application area · 3 min

The challenge target counts passes, time with the true mean altitude outside the band, the mean at the end and xenon. Predictive maintenance with one-burn reboosts and the uncertainty-aware planner meets it with 5.96 g in 10 burns, against 6.04 g in 12 burns with two-burn reboosts.

(a) Using case 3, say what the one-burn plan does to the orbit that the score cannot see. *Checkpoint:* the apogee rises 325 m above the band.

(b) An Earth-observation customer needs the altitude inside the band everywhere around the orbit. Name the quantities you would add to the target and where the app reports them. *Checkpoint:* the mean perigee and apogee.

8.12 Review the agent · 5 min

An AI assistant is asked to automate station keeping and replies as follows. “The mean altitude lags the orbit by half a revolution, so a planner on it is always late. The navigation fix gives the altitude every 10 s with only 5 m of noise. Close a PI loop from the measured altitude to the along-track thrust with a setpoint of 500 km, and let the thruster rest while the altitude is within the ±250 m band. The loop will hold the altitude with far less delay than any planner.” Find the flaw in the proposal and say how it could be checked.