

Guidance, navigation and control: know, decide, act

Close the loop: know, compare, correct

Starting question — How do guidance, navigation and control turn a mission objective into stable vehicle motion?

Intuition. Navigation estimates the state, guidance defines where the vehicle should go, and control acts on the difference. Confusing these roles hides failures in sensors, estimation or actuators.

- Explain the governing physical idea before calculating.
- Name every symbol and unit used in the key relation.
- Check the result with an independent inverse, bound or order-of-magnitude test.



Guidance and navigation are not isolated functions: they depend on a robust local network, relays, timing and autonomy adapted to the Earth-Mars delay.



Conceptual visualisation of an external measurement used to update navigation. The teaching principle is to combine sensors with different error modes so that the state estimate remains observable when one source temporarily degrades.

GNC—guidance, navigation and control—contains three distinct jobs. Navigation estimates the vehicle state. Guidance decides what state or trajectory should be pursued. Control commands actuators to reduce the error. Keeping those jobs separate makes both design and fault diagnosis clearer.

1. Separate navigation, guidance and control

Navigation estimates position, velocity, attitude and sensor biases. **Guidance** generates a desired trajectory or state. **Control** converts error into actuator commands. If navigation is wrong, a perfectly functioning controller can drive the spacecraft away from the real target while believing it is correcting an error.

Guidance, navigation and control form one loop but answer three different questions. Navigation estimates the actual state—position, velocity, attitude and often sensor biases. Guidance converts mission intent into a trajectory or attitude reference. Control commands actuators to reduce the difference between estimated state and reference. That separation is diagnostic: a spacecraft can track a bad reference perfectly, meaning control may be healthy while guidance is wrong. If the reference is correct but the state estimate drifts, navigation is the suspect. An explainable architecture preserves these boundaries in telemetry, fault messages and degraded modes so the crew knows which part of the chain has lost credibility.

2. Inertial sensors propagate motion but drift

An IMU combines gyroscopes and accelerometers. Integrating their measurements propagates attitude, velocity and position, but sensor bias accumulates. A constant acceleration bias $\mathbf{b} = 10^{-4} \text{ m/s}^2$ over $t = 1,000 \text{ s}$ creates roughly $\Delta \mathbf{v} = \mathbf{b} t = 0.1 \text{ m/s}$ of velocity error. Position error grows even more strongly. Inertial navigation therefore needs external updates.

An inertial measurement unit senses specific force and angular rate. Velocity, position and attitude are obtained by integrating those measurements over time, which means sensor bias is integrated as well. An accelerometer error of only 100 micro-g—about 0.000981 m/s^2 —would accumulate roughly 0.85 m/s of velocity error after 15 minutes in a deliberately simple one-dimensional calculation. Real navigation adds rotations, gravity models and filtering, but the lesson remains: an

IMU is excellent for short-term continuity, not for providing absolute truth by itself for days. External measurements must periodically constrain the estimate, and their uncertainty must also be represented.

Calculation laboratory — certified formula reasoning

This laboratory adds the zero-prerequisite reasoning chain used by the certified French Academy: read the formula, define every symbol and unit, calculate step by step, then verify the result independently.

Guidance, navigation and control: canonical mini-lessons

These mini-lessons form the module's canonical quantitative reference. Numerical cases are teaching examples and must remain distinct from mission data.

Tracking error

$$e = r - y$$

1 — Concrete question

What does “ $e = r - y$ ” compute in the context of “Tracking error”?

2 — Intuition without symbols

The controller compares commanded state with estimated or measured state; their difference is the error to reduce.

3 — Quantities

e : error; r : reference or command; y : measurement or estimate.

4 — Formula

$$e = r - y$$

5 — Read aloud

Read “ $e = r - y$ ” by naming every operation explicitly.

6 — Symbols and meaning

e : error; r : reference or command; y : measurement or estimate.

7 — Pronunciation

Say subscripts, exponents, Greek letters and operations explicitly so another person can reconstruct the relation without seeing the page.

8 — Units

Same unit for e , r and y , e.g. degrees.

9 — Convention

Keep the same frame, sign convention, time basis and system boundary throughout the calculation.

10 — Why this operation

Subtracting observed state from command gives the signed offset to correct.

11 — Assumptions

Same frame, sign convention and epoch for r and y .

12 — Unit check

Same unit for e , r and y , e.g. degrees. Verify that units reduce to the announced output quantity.

13 — Numerical case

With $r = 10^\circ$ and $y = 8^\circ$, $e = 10 - 8 = +2^\circ$.

14 — Why the calculation works

Subtracting observed state from command gives the signed offset to correct.

15 — Algebraic check

$y + e$ must recover r .

16 — Mental estimate

10 minus 8 immediately gives 2° .

17 — Interpretation

The sign indicates on which side of command the estimated state lies.

18 — What the result does not prove

This result proves neither universal validity of the model nor quality of the inputs; it applies only to the stated scenario and assumptions.

19 — Sensitivity

A measurement closer to command reduces $|e|$; measurement error can instead shift e artificially.

20 — Guided and autonomous exercises

Guided exercise. $r = 15^\circ$ and $y = 11^\circ$.

Detailed guided correction — open after trying

$$e = 15 - 11 = +4^\circ.$$

Autonomous exercise. $r = 5^\circ$ and $y = 7^\circ$.

Autonomous correction — open after trying

$$e = 5 - 7 = -2^\circ; \text{ sign reverses because estimate exceeds command.}$$

21 — Mission decision

Verify sign convention before closing a control loop.

Proportional control command

$$u = K \times e$$

1 — Concrete question

What does “ $u = K \times e$ ” compute in the context of “Proportional control command”?

2 — Intuition without symbols

A proportional command reacts to error by applying a gain: larger error, larger correction.

3 — Quantities

u : command; K : proportional gain; e : error.

4 — Formula

$$u = K \times e$$

5 — Read aloud

Read “ $u = K \times e$ ” by naming every operation explicitly.

6 — Symbols and meaning

u : command; K : proportional gain; e : error.

7 — Pronunciation

Say subscripts, exponents, Greek letters and operations explicitly so another person can reconstruct the relation without seeing the page.

8 — Units

Unit of u = unit of $K \times$ unit of e .

9 — Convention

Keep the same frame, sign convention, time basis and system boundary throughout the calculation.

10 — Why this operation

Multiplication represents how strongly command responds to error.

11 — Assumptions

Regime where local proportional control is relevant and unsaturated.

12 — Unit check

Unit of u = unit of $K \times$ unit of e . Verify that units reduce to the announced output quantity.

13 — Numerical case

With $e = 2^\circ$ and $K = 0.5$ command unit/ $^\circ$, $u = 0.5 \times 2 = 1.0$ command unit.

14 — Why the calculation works

Multiplication represents how strongly command responds to error.

15 — Algebraic check

u/K must recover $e = 2^\circ$.

16 — Mental estimate

A gain of one half applied to two error units gives one command unit.

17 — Interpretation

Too little K corrects slowly; too much K can encourage oscillation or saturation.

18 — What the result does not prove

This result proves neither universal validity of the model nor quality of the inputs; it applies only to the stated scenario and assumptions.

19 — Sensitivity

At fixed e, u varies linearly with K until system limits.

20 — Guided and autonomous exercises

Guided exercise. $K = 0.8$ command unit/ $^{\circ}$ and $e = 1.5^{\circ}$.

Detailed guided correction — open after trying

$u = 0.8 \times 1.5 = 1.20$ unit.

Autonomous exercise. $K = 0.3$ command unit/ $^{\circ}$ and $e = -4^{\circ}$.

Autonomous correction — open after trying

$u = 0.3 \times (-4) = -1.2$ unit; sign follows error.

21 — Mission decision

Choose K using real stability, saturation and dynamics, not this calculation alone.

Accelerometer-bias accumulation

$$\Delta v_{\text{bias}} \approx b_a \times t$$

1 — Concrete question

What does “ $\Delta v_{\text{bias}} \approx b_a \times t$ ” compute in the context of “Accelerometer-bias accumulation”?

2 — Intuition without symbols

A small persistent acceleration bias integrates over time into measurable velocity error.

3 — Quantities

Δv_{bias} : velocity error; b_a : acceleration bias; t : duration.

4 — Formula

$$\Delta v_{\text{bias}} \approx b_a \times t$$

5 — Read aloud

Read “ $\Delta v_{\text{bias}} \approx b_a \times t$ ” by naming every operation explicitly.

6 — Symbols and meaning

Δv_{bias} : velocity error; b_a : acceleration bias; t : duration.

7 — Pronunciation

Say subscripts, exponents, Greek letters and operations explicitly so another person can reconstruct the relation without seeing the page.

8 — Units

Δv in m/s; b_a in m/s²; t in s.

9 — Convention

Keep the same frame, sign convention, time basis and system boundary throughout the calculation.

10 — Why this operation

Constant acceleration integrated over t produces velocity change equal to acceleration times time.

11 — Assumptions

Bias treated as constant over the interval.

12 — Unit check

Δv in m/s; b_a in m/s²; t in s. Verify that units reduce to the announced output quantity.

13 — Numerical case

With $b_a = 1 \times 10^{-4} \text{ m/s}^2$ for 1,000 s, $\Delta v_{\text{bias}} = 0.10 \text{ m/s}$.

14 — Why the calculation works

Constant acceleration integrated over t produces velocity change equal to acceleration times time.

15 — Algebraic check

$\Delta v/t$ must recover $1 \times 10^{-4} \text{ m/s}^2$.

16 — Mental estimate

One ten-thousandth m/s^2 for one thousand seconds gives one tenth m/s .

17 — Interpretation

This shows why bias estimation and state updates are essential to autonomous navigation.

18 — What the result does not prove

This result proves neither universal validity of the model nor quality of the inputs; it applies only to the stated scenario and assumptions.

19 — Sensitivity

Doubling time doubles velocity error if bias stays constant.

20 — Guided and autonomous exercises

Guided exercise. $b_a = 2 \times 10^{-4} \text{ m/s}^2$ for 600 s.

Detailed guided correction — open after trying

$$\Delta v = 2 \times 10^{-4} \times 600 = 0.12 \text{ m/s}.$$

Autonomous exercise. $b_a = 5 \times 10^{-5} \text{ m/s}^2$ for 2,000 s.

Autonomous correction — open after trying

$$\Delta v = 5 \times 10^{-5} \times 2,000 = 0.10 \text{ m/s}.$$

21 — Mission decision

Set update frequency and error budget before bias consumes navigation margin.

Small-angle lateral error

$$x \approx R \times \theta$$

1 — Concrete question

What does " $x \approx R \times \theta$ " compute in the context of "Small-angle lateral error"?

2 — Intuition without symbols

A small angular error becomes lateral error that grows with range to target.

3 — Quantities

x : lateral error; R : range; θ : angular error in radians.

4 — Formula

$$x \approx R \times \theta$$

5 — Read aloud

Read " $x \approx R \times \theta$ " by naming every operation explicitly.

6 — Symbols and meaning

x : lateral error; R : range; θ : angular error in radians.

7 — Pronunciation

Say subscripts, exponents, Greek letters and operations explicitly so another person can reconstruct the relation without seeing the page.

8 — Units

x and R in the same length unit; θ in radians.

9 — Convention

Keep the same frame, sign convention, time basis and system boundary throughout the calculation.

10 — Why this operation

For a small angle, arc length and tangent are approximately $R\theta$.

11 — Assumptions

Small angle expressed in radians and range clearly defined.

12 — Unit check

x and R in the same length unit; θ in radians. Verify that units reduce to the announced output quantity.

13 — Numerical case

At $R = 8,000$ m and $\theta = 0.15^\circ \times \pi/180 = 0.002618$ rad, $x \approx 8,000 \times 0.002618 = 20.94$ m.

14 — Why the calculation works

For a small angle, arc length and tangent are approximately $R\theta$.

15 — Algebraic check

x/R must recover about 0.002618 rad.

16 — Mental estimate

At 8 km, one milliradian is about 8 m; 2.6 milliradians therefore give a little over 20 m.

17 — Interpretation

The same pointing error becomes more costly as target range increases.

18 — What the result does not prove

This result proves neither universal validity of the model nor quality of the inputs; it applies only to the stated scenario and assumptions.

19 — Sensitivity

x varies linearly with R and θ in the small-angle regime.

20 — Guided and autonomous exercises

Guided exercise. $R = 5,000$ m and $\theta = 0.10^\circ = 0.001745$ rad.

Detailed guided correction — open after trying

$x \approx 5,000 \times 0.001745 = 8.73$ m.

Autonomous exercise. $R = 12,000$ m and $\theta = 0.05^\circ = 0.000873$ rad.

Autonomous correction — open after trying

$x \approx 12,000 \times 0.000873 = 10.47$ m.

21 — Mission decision

Translate lateral-position requirement into pointing and navigation requirements.

Mission reasoning lab — diagnose a guidance, navigation and control loop

Scenario. A lander is commanded to hold a 10-degree attitude. Navigation estimates 8 degrees, so the control error is +2 degrees. A few seconds later the estimate moves to 10.5 degrees while the command remains unchanged, so the error becomes -0.5 degree. This sign change is the essence of negative feedback: correction must reverse once the estimate crosses the command.

1. Keep guidance, navigation and control separate

Guidance decides the target state or trajectory. Navigation estimates the actual state. Control converts their difference into actuator commands. If a vehicle goes to the wrong place, engineers must determine which layer failed. A perfect controller faithfully following a biased navigation solution can still drive the vehicle away from the true target.

2. Frames and timestamps are part of the equation

Subtracting a body-frame angle from an inertial-frame angle is meaningless until transformed into a common frame. The same is true when one state is 200 milliseconds old and the other is current in a fast dynamic phase. The algebraic relation $\text{error} = \text{command} - \text{estimate}$ is simple; the engineering validity of its inputs is not.

3. Use inverse checks

With command = 10° and error = $+2^\circ$, estimated state must be 8° . With estimate = 8° and error = $+2^\circ$, command must be 10° . These inverse relations are trivial enough to do mentally and are excellent at detecting sign errors in software interfaces.

4. Distinguish estimation uncertainty from control error

An estimated error of zero does not prove the true state is correct. If navigation uncertainty is $\pm 3^\circ$, a displayed 0.1° error can coexist with a much larger true error. Controllers and mission rules therefore need uncertainty bounds, sensor health and mode logic, not only a point estimate.

5. Actuator authority and saturation

Suppose the controller asks for a torque greater than the actuators can produce. The numerical command may grow while the physical response saturates. Integrators can wind up and recovery can become slow or unstable. A robust design detects saturation and changes control strategy rather than assuming every requested command is physically achievable.

6. Delay and stability

Feedback always contains some delay from sensing, filtering, computation and actuation. Excessive delay can turn a stabilizing correction into an oscillation because the controller acts on an outdated state. A Mars surface robot also faces long Earth communication delay, so local control must remain autonomous even when high-level guidance comes from Earth.

Decision check

Before trusting a GNC result, ask: Are command and estimate in the same frame and unit? Are timestamps aligned? Is uncertainty acceptable? Is the sign convention documented? Are actuators within authority? Is delay included in stability analysis? A loop that answers all six is far more credible than one that only shows a small error value.

Fault cases in a closed control loop

Wrong-sign feedback. If the controller adds rather than removes error, even a small disturbance grows. A sign mistake can come from axis definitions, actuator polarity or coordinate transforms. Hardware-in-the-loop testing should deliberately inject small known disturbances and verify that commanded response drives the state back toward the target.

Sensor disagreement. Redundant sensors do not automatically create a correct estimate. Common-mode bias, shared software or identical environmental sensitivity can make several sensors agree and still be wrong. Navigation logic needs independence arguments, innovation monitoring and rules for rejecting or reweighting suspect measurements.

Mode transitions. Launch, cruise, entry and landing can use different estimators and controllers. The transition itself is a hazard: state definitions, gains, actuator sets and reference frames can change. A safe design defines entry criteria, handover logic, rollback conditions and post-transition checks.

Operational observability. Mission control needs more than a green "GNC nominal" light. It needs enough telemetry to distinguish sensor failure, estimator divergence, actuator saturation and guidance-command error. Good observability shortens diagnosis time and prevents the wrong recovery action.

Zero-prerequisite concepts

These concept cards define the vocabulary that a first-time learner must master before using it operationally.

guidance

Definition. Guidance determines the desired trajectory or commanded state needed to reach mission objectives.

Example. An entry guidance law may command bank angle or attitude targets to manage range and energy.

Pitfall. Guidance does not by itself measure the vehicle state.

Mental check. If the mission target changes, guidance commands may change even when navigation estimates remain identical.

Guided exercise — guidance

Situation to recognize. Guidance determines the desired trajectory or commanded state needed to reach mission objectives.

Check requested. If the mission target changes, guidance commands may change even when navigation estimates remain identical.

Error to reject. Guidance does not by itself measure the vehicle state.

Reasoned solution

Precise meaning

Guidance determines the desired trajectory or commanded state needed to reach mission objectives.

Case test

If the mission target changes, guidance commands may change even when navigation estimates remain identical.

Excluded pitfall

Guidance does not by itself measure the vehicle state.

Operational consequence

Use this check before accepting a result in mission design: If the mission target changes, guidance commands may change even when navigation estimates remain identical.

Quantification

guidance: use the unit or dimension defined by the physical quantity; if the concept is qualitative, do not invent a numerical unit.

Verification

guidance: compare the conclusion with the mental check and the stated pitfall.

navigation

Definition. Navigation estimates the vehicle's current state—such as position, velocity and attitude—from sensors and models.

Example. A spacecraft fuses inertial measurements, star tracking and radio data to estimate where it is.

Pitfall. Navigation is an estimate and always carries uncertainty.

Mental check. A sensor bias can make navigation wrong even while the controller follows its estimate perfectly.

Guided exercise — navigation

Situation to recognize. Navigation estimates the vehicle's current state—such as position, velocity and attitude—from sensors and models.

Check requested. A sensor bias can make navigation wrong even while the controller follows its estimate perfectly.

Error to reject. Navigation is an estimate and always carries uncertainty.

Reasoned solution

Precise meaning

Navigation estimates the vehicle's current state—such as position, velocity and attitude—from sensors and models.

Case test

A sensor bias can make navigation wrong even while the controller follows its estimate perfectly.

Excluded pitfall

Navigation is an estimate and always carries uncertainty.

Operational consequence

Use this check before accepting a result in mission design: A sensor bias can make navigation wrong even while the controller follows its estimate perfectly.

Quantification

navigation: use the unit or dimension defined by the physical quantity; if the concept is qualitative, do not invent a numerical unit.

Verification

navigation: compare the conclusion with the mental check and the stated pitfall.

control

Definition. Control converts guidance error into actuator commands that change vehicle motion.

Example. Thrusters, reaction wheels or aerodynamic surfaces respond to controller outputs.

Pitfall. A controller cannot correct a state it does not observe or estimate adequately.

Mental check. If actuator authority is saturated, reducing the numerical error in software alone cannot recover control.

Guided exercise — control

Situation to recognize. Control converts guidance error into actuator commands that change vehicle motion.

Check requested. If actuator authority is saturated, reducing the numerical error in software alone cannot recover control.

Error to reject. A controller cannot correct a state it does not observe or estimate adequately.

Reasoned solution

Precise meaning

Control converts guidance error into actuator commands that change vehicle motion.

Case test

If actuator authority is saturated, reducing the numerical error in software alone cannot recover control.

Excluded pitfall

A controller cannot correct a state it does not observe or estimate adequately.

Operational consequence

Use this check before accepting a result in mission design: If actuator authority is saturated, reducing the numerical error in software alone cannot recover control.

Quantification

control: use the unit or dimension defined by the physical quantity; if the concept is qualitative, do not invent a numerical unit.

Verification

control: compare the conclusion with the mental check and the stated pitfall.

feedback loop

Definition. A feedback loop repeatedly measures or estimates the state, compares it with a target and applies corrective action.

Example. Attitude control closes a loop many times per second.

Pitfall. Feedback can become unstable if delay, sign or gain is wrong.

Mental check. A stable negative-feedback loop should normally reduce a small disturbance rather than amplify it.

Guided exercise — feedback loop

Situation to recognize. A feedback loop repeatedly measures or estimates the state, compares it with a target and applies corrective action.

Check requested. A stable negative-feedback loop should normally reduce a small disturbance rather than amplify it.

Error to reject. Feedback can become unstable if delay, sign or gain is wrong.

Reasoned solution

Precise meaning

A feedback loop repeatedly measures or estimates the state, compares it with a target and applies corrective action.

Case test

A stable negative-feedback loop should normally reduce a small disturbance rather than amplify it.

Excluded pitfall

Feedback can become unstable if delay, sign or gain is wrong.

Operational consequence

Use this check before accepting a result in mission design: A stable negative-feedback loop should normally reduce a small disturbance rather than amplify it.

Quantification

feedback loop: use the unit or dimension defined by the physical quantity; if the concept is qualitative, do not invent a numerical unit.

Verification

feedback loop: compare the conclusion with the mental check and the stated pitfall.

3. Star trackers provide an absolute attitude reference

A star tracker matches an observed star field to a catalogue and estimates orientation. It can be extremely accurate yet unavailable near bright bodies or in prohibited Sun angles. Gyroscopes bridge gaps between stellar solutions. The architecture combines high-rate propagation with intermittent absolute reference.

A star tracker recognises star patterns and provides a precise absolute attitude reference. Its availability still depends on field of view, stray light, temperature, angular rate and catalogue quality. Sun intrusion or occultation can remove a valid solution without any hardware failure. Software must therefore distinguish 'no solution now' from 'sensor dead'. During an outage, the IMU propagates attitude with growing uncertainty. When star-tracker data returns, the first measurements should be checked before full reintegration. That controlled transition prevents a single bad reacquisition from creating a large attitude command and makes recovery an evidence-based process rather than an automatic switch.

4. State estimation combines a model with uncertain measurements

An estimator predicts the state from dynamics and then corrects that prediction with measurements and uncertainty models. The measurement-minus-prediction residual is valuable for fault detection. An overconfident model can reject valid sensor evidence; an overpermissive model can follow noise.

Exercise A — intuitive weighting

Two sensors report 100 and 104, but the first is known to be much more precise. Is 102 automatically the best estimate?

No. A weighted estimate gives more influence to the measurement with lower uncertainty. A simple average silently assumes equal quality.

A state estimator combines a dynamic prediction with uncertain measurements. The important idea is not the filter's name but the balance between model and evidence. A noisy measurement should move the estimate only a little; a precise, consistent measurement can correct it strongly. The innovation or residual compares the expected measurement with what was received. An improbable residual can indicate a bad sensor, an incorrect model or an unmodelled manoeuvre. Covariance represents estimated uncertainty and correlation. GNC should therefore publish both state and confidence, because the same position estimate might be good enough for antenna pointing and completely inadequate for terminal descent.

5. Observability asks whether a state can actually be known

A system can carry many sensors yet be unable to distinguish two candidate states if they produce the same measurements. A change in geometry, attitude or sensing technology may make the state observable. Observability belongs to the combination of dynamics and measurements rather than to sensor count alone.

Observability asks whether available measurements can actually separate the states being estimated. Many sensors do not automatically create independent information. Repeating a range-only measurement, for example, can leave several positions geometrically compatible if there is no angular information or helpful vehicle motion. A second line of sight or a

change in geometry can make the state observable. Before adding a sensor, the engineer asks which previously ambiguous state it constrains. This prevents superficial redundancy in which several instruments share the same blind direction, reference or error model and therefore fail together when the geometry becomes unfavourable.

6. Guidance must request a reachable reference

“Go to the safe site” is incomplete. Guidance considers vehicle dynamics, thrust limits, forbidden corridors and remaining reserve. An impossible reference simply drives the controller into saturation. Good guidance selects a physically reachable target inside the current capability envelope.

Guidance must request a reference that the vehicle can physically reach. A mathematically smooth path may still demand acceleration, angular rate or propellant beyond actuator capability. Guidance therefore includes thrust limits, allowable attitude, obstacles, propellant reserve and navigation margin. During descent, a late divert can saturate engines or force an attitude that degrades sensors. A sound design first computes a reachable set and then selects a reference inside it. Control should not be expected to make an impossible command possible. This is one of the most important boundaries between trajectory planning and feedback control.

7. Control closes the error loop

In a basic representation, $\mathbf{e} = \mathbf{r} - \mathbf{y}$ compares reference $\mathbf{r}$ with measured output $\mathbf{y}$. A proportional command might be $\mathbf{u} = \mathbf{K} \mathbf{e}$. $\mathbf{K}$ is controller gain and $\mathbf{u}$ actuator command. Larger gain is not automatically better; delay, noise, actuator limits and unmodelled dynamics can make an aggressive loop oscillate.

Exercise B — proportional command

An attitude error is 2 degrees and a teaching controller uses $K = 0.5$ command units per degree. What raw command results?

$u = K e = 0.5 \times 2 = 1$ command unit. Real design must then check sign, units, saturation and closed-loop dynamics.

Control closes the loop by turning error into command. Even a simple proportional example, $u = K \times e$, shows the trade: too little gain K gives slow correction; too much can excite delays, structural flexibility, sensor noise or actuator saturation. Real vehicles add integral and derivative terms, filters, limiters and actuator dynamics. Stability therefore means more than ‘the error eventually reaches zero’. Engineers also check overshoot, settling time, robustness to uncertain parameters and behaviour under noisy measurements. A survival mode may deliberately accept poorer pointing accuracy in exchange for wider stability margin and simpler logic.

8. Actuators saturate and carry resources

Reaction wheels can accumulate momentum and reach speed limits. Thrusters consume propellant and have minimum impulse characteristics. The controller needs these constraints in its model. Demanding more from a saturated actuator does not create more control authority.

Actuators carry their own resources and limits. A reaction wheel stores angular momentum until saturation; a thruster consumes propellant and has a minimum impulse bit; a gimbal has travel and rate limits; a magnetorquer depends on a local magnetic field and is not a deep-space attitude solution. GNC must know these boundaries. As wheels approach saturation, momentum unloading transfers angular momentum through another actuator. Operations must schedule that activity so it does not corrupt a precision observation or manoeuvre. Control authority is therefore a budgeted mission resource, much like energy or propellant.

9. GNC fault management needs independent evidence

FDIR can use residuals, cross-sensor consistency, expected dynamics and health flags. Three identical sensors may still fail together through shared software, calibration or environment. Diversity across inertial, stellar, radio, optical and terrain-relative references can be more valuable than blind replication.

GNC fault management is more than majority voting among three sensors. Identical sensors can share software, power, catalogue data or environmental sensitivities and therefore fail from one common cause. Stronger FDIR compares evidence of different kinds: inertial, stellar, solar, radio and expected dynamics. It also checks time history. A sudden jump in one sensor while others and actuators remain coherent is different from a slow common drift that may indicate a bad model. Isolation must preserve enough information for the next safe mode; otherwise fault management can create the very loss of control it was meant to prevent.

10. Failure scenario: star tracker unavailable during a manoeuvre

The vehicle propagates attitude on gyroscopes while uncertainty grows. A proper degraded-mode requirement says how long that propagation remains acceptable, which activities are suspended and what alternate measurement can re-establish absolute reference. “Fly inertially” is not complete until its validity envelope is quantified.

During a manoeuvre without a star tracker, the IMU can preserve attitude for a limited period, but uncertainty grows. The spacecraft needs a predeclared accuracy requirement for each activity. Coarse solar pointing may tolerate several tenths of a degree while a precision burn may not. If estimated uncertainty crosses the manoeuvre threshold, the correct action can be to stop and reacquire an absolute reference rather than continue on dead reckoning. Training should include sensor recovery, consistency testing, controlled filter reintegration and proof that actuators did not saturate during the outage. Recovery is complete only when the state and its confidence again satisfy the mission requirement.

Guided case — an estimator that becomes overconfident

Imagine an attitude estimator whose covariance steadily shrinks because star-tracker measurements have been very stable. A stray-light condition then creates a small persistent bias that the filter model does not represent. If software has become overconfident, it may reject a correct inertial observation as ‘inconsistent’. The danger is not only sensor noise but a poor representation of uncertainty. Diagnosis therefore compares innovation, sensor context and independent references before isolating a channel.

One recovery strategy temporarily inflates uncertainty, reduces star-tracker weighting and uses Sun sensing or radio geometry to discriminate the hypotheses. The filter is not merely a precision algorithm; it is an arbiter of confidence. Robust missions monitor the statistical consistency of residuals over time rather than watching only the state estimate itself.

The exercise asks students to distinguish three failures: a bad measurement, an incorrect dynamic model and an unrealistically small covariance. All three can create the same symptom—a large residual—but require different responses. That distinction prevents the team from replacing the most visible sensor when the real problem may be in software or assumptions.

11. Mini-project: build a GNC evidence chain

1. Select attitude, altitude or position as the controlled quantity.
2. List sensors that observe it directly or indirectly.
3. Define the guidance reference.
4. Identify the actuator.
5. Add one saturation and one lost sensor.
6. Explain detection, isolation and the degraded state.

The final diagram should trace information all the way from measurement to physical effect.

12. Mission lab — detect inertial drift without creating a false alarm

A gyro bias of only 0.01 degrees per hour looks negligible. After 24 hours without an absolute update, a first-order attitude error scale is 0.24 degrees. The relation $\Delta\theta \approx \mathbf{b} \times \mathbf{t}$ uses $\Delta\theta$ for angle error, $\mathbf{b}$ for angular bias per unit time and $\mathbf{t}$ for elapsed time. A real navigation filter estimates bias and attitude together, but the simple multiplication shows why periodic absolute references matter.

Now assume the star tracker and IMU begin to disagree slowly. Good FDIR first checks geometry and context: is the star field partially blocked, is the tracker near a thermal or Sun-avoidance limit, did the IMU just experience a high-dynamics manoeuvre? A third independent reference—Sun sensing, radio geometry or another optical observation—may resolve the disagreement. Immediately declaring the IMU faulty could discard the only healthy source.

State estimates should carry confidence. An attitude solution with growing uncertainty may remain adequate for solar-array pointing while no longer meeting a precision manoeuvre requirement. Operational limits therefore belong to activities, not to a single universal “navigation valid” flag.

This idea extends to position and velocity. The same estimated state can be acceptable for a coarse communication pointing mode and unacceptable for terminal descent. GNC health must be judged against the current mission need.

13. Stability, saturation and controlled return to nominal

An actuator that saturates cannot deliver more authority simply because the controller asks harder. Integrating controllers can continue accumulating error while saturated, producing overshoot when authority returns. Anti-windup, reference limiting or mode switching are examples of remedies. The main lesson is that control laws live inside physical amplitude, rate, energy and thermal limits.

Sensor recovery also needs a transition. A star tracker that comes back after an outage can be compared against propagated attitude, checked for residual consistency and gradually accepted. An abrupt switch to a measurement with an unnoticed offset can create a command transient. Return-to-service criteria should state allowable residual, required observation duration and what happens if disagreement reappears.

For a design review, trace one complete off-nominal loop: sensor fault, detector evidence, confidence change, guidance restriction, actuator command, safe-state entry, new measurement, recovery decision and restored mission activity. If any link is described only as “software handles it,” the GNC case is incomplete.

Beginner vocabulary checkpoint

These terms are intentionally interactive: focus or tap a term for its definition, then return to the surrounding calculation or mission decision.

- guidance — Logic that determines where the vehicle should go and what trajectory or attitude it should command.
- navigation — Estimation of the vehicle state from sensors, models and external measurements.
- control — Actions that drive the estimated state toward the commanded state.
- command — Desired value supplied to a controller, such as attitude, speed or position.
- measurement — Observed sensor quantity before or after calibration and filtering.
- state estimate — Best current estimate of variables such as position, velocity and attitude.
- residual — Difference between an observed quantity and its predicted value; useful for fault detection.
- Kalman filter — Estimator that combines a dynamic model with noisy measurements using uncertainty information.
- sensor bias — Persistent measurement offset that can create systematic navigation error.
- noise — Random variation that obscures the underlying measured signal.
- actuator — Device that physically changes vehicle state, such as a thruster, reaction wheel or control surface.
- control law — Rule that maps state error into actuator commands.
- feedback — Use of measured or estimated response to adjust subsequent control action.
- open loop — Command sequence executed without using response feedback to correct the action.
- closed loop — Control architecture that repeatedly compares commanded and estimated states.
- deadband — Error interval inside which no corrective action is commanded.
- fault detection — Process used to identify abnormal sensor, actuator or software behaviour.
- redundancy — Use of independent alternatives so one failure does not automatically remove a critical function.

Sources and references

- NASA — State of the Art: Guidance, Navigation and Control
- NASA/JPL — Mars 2020 Landing and Terrain-Relative Navigation
- NASA SCA_N — Space Communications and Navigation

Verified primary supplement: NASA NTRS — State-of-the-Art Small Spacecraft Technology

Engineering studio — connect angular error to safety

At 8 km from a target, an angular error of 0.15° corresponds to $\theta = 0.15 \times \pi / 180 \approx 0.00262$ rad. For small angles, lateral error is approximately $x = R\theta = 8,000 \times 0.00262 \approx 21$ m. Here R is range in metres and θ is angle in radians. The calculation shows why an apparently tiny sensor error can become tens of metres on the ground.

The scenario then adds 200 ms of processing latency and limits one actuator to 80% of nominal authority. The student decides whether the error remains observable and recoverable before the next decision gate. At least two independent checks are required — for example filter innovation and consistency with an optical measurement — followed by a criterion that triggers a more conservative trajectory.

FR ↔ EN competency matrix

#	Français	English	Status
01	Erreur de suivi	Tracking error	PASS
02	Commande proportionnelle	Proportional control command	PASS
03	Accumulation d'un biais d'accéléromètre	Accelerometer-bias accumulation	PASS
04	Erreur transversale de petit angle	Small-angle lateral error	PASS