System dynamics, feedback and estimation

Start from a physical plant and its measured signals. Derive a feedback or estimation model and test which dynamical behaviour can be controlled or inferred.

Contents

1. **Estimating velocity from position with a model-based observer**
   An optical encoder returns position and the actuator drive supplies a commanded acceleration. Reconstruct an unmeasured velocity without differentiating successive noisy samples. Derive the observer, explain what a pair of chosen poles guarantees, and predict the tradeoff visible when the encoder carries periodic interference.
2. **Relative gains, loop pairing and nearly singular process measurements**
   Two heaters act on two temperature sensors. With the second heater fixed, a test of the first gives a well-determined positive response. Closing a loop around the other temperature can change both the size and sign of that apparent gain. Derive the relative gain array from the experiment one actually performs.
3. **Predicting a delayed process and testing a Smith controller**
   A heater is upstream of a temperature sensor. Its internal thermal response begins immediately while a material parcel reaches the sensor several seconds later. Build a feedback signal from two model outputs, derive which transfer function loses its delay, and then deliberately introduce a delay mismatch to see why a plausible prediction may fail.
4. **Sensitivity, bandwidth and the integral price of disturbance rejection**
   A regulator attenuates a slowly varying output load yet amplifies another disturbance at a higher frequency. Separate the paths by which reference, measurement noise and physical disturbance reach the recorded variable, then derive an integral identity that explains why a stable loop can redistribute sensitivity across frequency.
5. **Which sampling times reveal the velocity of an oscillator?**
   A sensor measures displacement while a controller needs velocity as well. Use the exact free vibration between two readings to recover the missing state, then quantify how a poorly chosen sample interval amplifies measurement noise.
6. **When does a stable operating point disappear?**

A shaft can hold a static torque against gravity up to a finite limit. Trace the equilibrium branches and distinguish local stability, loss of an equilibrium and the possibility of oscillation.

7. Can any controller meet this bandwidth and peak-gain specification?

Delay uses up phase at the crossover even when it changes no gain. An unstable pole near a right-half-plane zero imposes a separate peak-gain limit. Derive each restriction before tuning a controller.

8. Designing current, speed and position loops for a drive

An encoder measures shaft motion, while the power stage acts on voltage. Decide which disturbances each measurement can reject and how fast each loop can move before the unseen inner dynamics erode the position margin.

System dynamics, feedback and estimation · 1. Estimating velocity from position with a model-based observer

An optical encoder returns position and the actuator drive supplies a commanded acceleration. Reconstruct an unmeasured velocity without differentiating successive noisy samples. Derive the observer, explain what a pair of chosen poles guarantees, and predict the tradeoff visible when the encoder carries periodic interference.

Find out which initial conditions are distinguishable

Model a translating stage over its usable travel by $\dot{x}_1 = x_2,\ \dot{x}_2 = u$, with measured output $y = x_1 + n$. Here the input is an acceleration in the chosen position and time units. With zero noise and known input, two different initial positions give different initial readings. Two different initial velocities with the same position yield traces whose initial slopes differ. Formally the observability matrix stacks $C = (1, 0)$ above $CA = (0, 1)$, giving rank two. If only velocity were measured, a constant offset of initial position would remain invisible, whatever the command. Observability belongs to the chosen model and output together.

Run a copy of the plant and drive it with both the known command and a correction proportional to innovation $y - \hat{x}_1$:

$$\dot{\hat{x}}_1 = \hat{x}_2 + l_1(y - \hat{x}_1),$$

$$\dot{\hat{x}}_2 = u + l_2(y - \hat{x}_1).$$

Define error as true minus estimated state. For a perfect model and noise-free measurement the error matrix has rows $(-l_1, 1)$ and $(-l_2, 0)$. Expanding the determinant of $sI - (A - LC)$ gives $s^2 + l_1 s + l_2$. Matching the standard oscillator denominator sets $l_1 = 2\zeta\omega_o$ and $l_2 = \omega_o^2$. The units are inverse seconds and inverse seconds squared. A constant position innovation can therefore request corrections to both estimated velocity and acceleration. Stable poles guarantee decay of initial error for the ideal model. They do not specify how large the measurement error becomes when injected continually.

Follow one sensor tone all the way to the velocity estimate

Set the true state and command to zero and solve the two linear equations in the Laplace domain with zero initial condition. Eliminating estimated position gives the transfer of position noise to estimated velocity:

$$\frac{\hat{V}(s)}{N(s)} = \frac{l_2 s}{s^2 + l_1 s + l_2}.$$

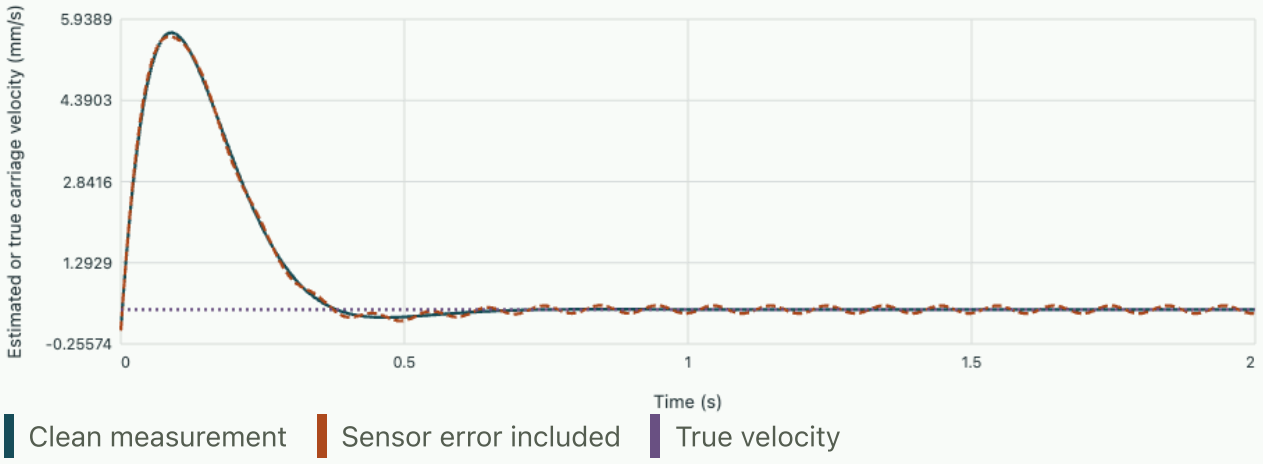
For slow variation it behaves like a differentiator, so a constant position offset ultimately produces zero spurious velocity. At high frequency the gain falls as l_2/ω, whereas differentiating raw data multiplies its amplitude by frequency. Noise near the observer bandwidth can nevertheless be strongly visible. With observer frequency 2 Hz, damping ratio 0.75 and a 10 Hz sensor oscillation of amplitude 0.03 mm,

the predicted velocity ripple is about 0.075 mm/s. The stage moves at 0.40 mm/s in the example. Compare the late-time difference between clean and noisy traces with the predicted peak ripple. The initial position estimate is deliberately displaced to reveal transient convergence.

Open the live experiment

Try the calculation

Observer natural frequency (Hz)	**2**
Sinusoidal sensor error (mm)	**0.03**
Sensor error frequency (Hz)	**10**



The predicted steady velocity ripple amplitude from the sensor error is 0.075 mm/s. Initial true position is 1 mm; both estimated initial states are zero.

Background · What if successive position errors are correlated?

Estimate velocity as the difference between consecutive positions divided by the sample interval h. Suppose each position error has variance σ^2 and the correlation between successive errors is ρ. The variance of the difference contains the covariance twice with a negative sign.

$$\operatorname{Var}\left(\frac{e_k - e_{k-1}}{h}\right) = \frac{2\sigma^2(1-\rho)}{h^2}.$$

A constant sensor bias cancels exactly from this difference, although the reconstructed position can remain wrong. Independent 0.10 mm position errors sampled 0.01 s apart give a 14.1 mm/s velocity standard deviation. With the same marginal variance and correlation 0.9 it falls to 4.47 mm/s. A sensor data sheet reporting only RMS position noise leaves this prediction underdetermined.

Use the innovation as a diagnostic signal

If the true acceleration exceeds the commanded value by constant d, the steady error equations yield

$$e_1 = d/l_2,$$

$$e_2 = l_1 d/l_2.$$

Raising the observer frequency reduces these model-error terms while passing more encoder noise at a fixed sufficiently high frequency. Adding a slowly evolving bias state is another modelling choice; its observability and process-noise assumptions need to be checked again. A near-zero innovation in a high-gain observer alone is weak evidence of correct velocity, since the algorithm was designed to force estimated position toward the sensor.

On the bench, estimate the encoder noise spectrum with the stage held stationary and the illumination and sampling electronics operating normally. During motion, compare estimated speed with an independently timed travel between two reference marks. An encoder fitted only by its own residual offers no independent absolute scale. If a motion command changes numerically from metres to millimetres, transform acceleration, position and initial estimates together. Pole locations have units of inverse time and do not depend on this uniform length conversion.

Check whether an apparent velocity ripple aliases from above the sampling band

Suppose the encoder is sampled at 100 Hz and its electronics pick up a 90 Hz sinusoid. At those sampling instants the measured sequence is indistinguishable from a ten-hertz sinusoid with suitable phase. A downstream digital observer cannot determine which analogue frequency generated the samples. Repeating with a different acquisition clock moves an alias while leaving a genuine ten-hertz process feature at its physical frequency. An analogue anti-alias filter can attenuate the high-frequency signal before sampling.

Compare this experiment with the frequency knob in the displayed model. The browser directly synthesises a known tone at the integration rate and therefore tests the observer frequency response in isolation. A measured disagreement on a real sampled encoder can reveal an upstream acquisition effect even when the observer discretisation is accurate.

Tutorial · Identify a constant offset and an acceleration mismatch

An encoder reads 0.20 mm high at rest. Separately, an unmodelled 0.20 mm/s squared acceleration acts on an otherwise perfect stage. Find steady errors for the two cases at the default observer setting.

▼ **Solve the steady error equations for each input**

A sensor offset eventually places estimated position 0.20 mm above truth while the estimated velocity returns to zero. For the acceleration mismatch, $l_2 = 157.91$ and $l_1 = 18.85$ in the appropriate second-

based units. The true position exceeds the estimate by 0.00127 mm and true velocity exceeds its estimate by 0.0239 mm/s. The persistent residual in the second case is tiny despite a clear velocity bias. Repeat the test while changing bandwidth to distinguish a bias in the sensor from a mismatch in acceleration.

Tutorial · Compare with numerical differentiation

With 1 ms sample separation, independent position errors have standard deviation 0.01 mm per reading. What standard deviation does a backward difference produce for velocity? Explain how this differs from the deterministic 10 Hz test.

▼ **Propagate the variance of the difference**

Subtracting independent zero-mean errors adds their variances. Divide the resulting standard deviation by the time step to obtain

$$\sqrt{2}(0.01)/0.001 = 14.14\,\mathrm{mm/s}.$$

A 10 Hz interference tone has correlated neighbouring samples and a fixed phase evolution. Its output is predicted from sinusoidal frequency response, not from independence. Averaging a longer interval reduces random position-difference noise but measures a longer-time mean speed and responds more slowly to acceleration.

Tutorial · Check discrete implementation before increasing gain

A software observer updates by explicit Euler every 10 ms. If the continuous error poles are both minus 300 per second, what happens to an initial error? How could the implementation preserve the intended decay?

▼ **Map a continuous pole through the actual update**

The scalar Euler factor is

$$1 + h\lambda = 1 - 3 = -2;$$

its magnitude exceeds one, so errors grow while changing sign. Exact discretisation would map the pole to $e^{-3} = 0.0498$. A suitably faster sample interval, a correctly discretised observer, or a discrete observer designed directly for the sampled measurement can restore stability. These choices still depend on the actual measurement rate. Interpolating old encoder samples at a faster software clock does not create independent new position information.

Caltech · Åström and Murray, Output Feedback, observability and state estimation

System dynamics, feedback and estimation · 2. Relative gains, loop pairing and nearly singular process measurements

Two heaters act on two temperature sensors. With the second heater fixed, a test of the first gives a well-determined positive response. Closing a loop around the other temperature can change both the size and sign of that apparent gain. Derive the relative gain array from the experiment one actually performs.

Compare two physically different partial derivatives

Linearise a two-input, two-output process about an equilibrium, using deviations from the original input and output settings. Write its steady gain matrix as K, with entry $k_{ij} = \partial y_i / \partial u_j$ with the other input held fixed. If instead the second output is held fixed by feedback, the other input must move too. Eliminate it from the second row:

$$0 = k_{21}\,du_1 + k_{22}\,du_2,$$

$$\left.\frac{dy_1}{du_1}\right|_{y_2} = k_{11} - \frac{k_{12}k_{21}}{k_{22}} = \frac{\det K}{k_{22}}.$$

The relative gain for this pairing is the open-other-input gain divided by the gain with the other output held constant. Inverting the gain matrix gives the general expression $\Lambda_{ij} = k_{ij}(K^{-1})_{ji}$, an element-by-element product with the transposed inverse. It requires a nonsingular square matrix; the alternative ratio interpretation requires that the selected constrained experiment also be well defined. Each row and column sums to one. A quick check uses the diagonal of $KK^{-1} = I$. A negative element signals a gain reversal between the two experiments. A very large element indicates considerable cancellation in the constrained response.

For matrices $K = \left(\begin{smallmatrix} 2 & 1 \\ 1 & 2 \end{smallmatrix}\right)$, the diagonal relative gains are 4/3 and the off-diagonal entries are minus one third. A unit increase of the first input holding the second output fixed requires reducing the second input by one half. The first output then rises by 1.5, compared with 2 when the other input is held fixed. Off-diagonal pairing would confront a reversal when the other loop closes. A static array can guide selection, but faster dynamics, delays, actuator bounds and output priorities still matter.

Show which numerical conclusions depend on chosen units

Let individual output units be rescaled by a diagonal matrix D_y and input units by D_u. The new gain matrix is $D_y K D_u^{-1}$ and its inverse transforms in the opposite order. The scale factors cancel separately in each product defining a relative gain. Reporting a temperature deviation in millikelvin instead of kelvin does not change the array. Singular values, by contrast, do depend on units; meaningful comparisons require physical normalisation such as allowed input ranges and tolerated output deviations.

For $K = \left(\begin{smallmatrix} 1 & 0.9 \\ 1.1 & 1 \end{smallmatrix}\right)$ the determinant is 0.01. The diagonal entries of the array reach 100 and the off-diagonal entries minus 99. Nominally holding the second output constant requires $du_2 = -1.1 du_1$, leaving only 0.01 output units per unit of first input. A requested unit output therefore asks for input

changes of 100 and minus 110 in this local model. Saturation or departure from the linear neighbourhood would usually intervene first. A nearly singular response indicates that the two actuators push in almost the same output direction.

A test programme distinguishes cancellation from sensor drift

Perturb each input on both sides of the reference point with the other held fixed, allow slow transients to settle, and fit slopes including an offset for drift. Repeat with both inputs coordinated to hold one output near its initial value. The small difference of large responses may have uncertainty as large as its point estimate. Record confidence intervals for the matrix elements and propagate them through the determinant before prescribing a controller pairing. Independent tests on a second operating point reveal whether the conclusion survives gain changes. A temperature sensor substituted into both output channels can accidentally make an otherwise invertible plant look rank deficient.

Compare the time-domain inverse-response study with a steady relative gain. A right-half-plane zero can constrain a pairing that looks attractive at zero frequency.

Make the nearly invisible constrained response visible

The first curve holds the second input constant; the second holds the second output constant. Both pass through the reference point because the model uses deviations about a common equilibrium. With the nominal gains the latter slope is only one hundredth of the former. Change the lower-right coefficient from one to 0.99 and then to 0.98. The constrained slope passes through zero and reverses sign. At the singular value a finite relative gain is undefined, while the physical experiment remains perfectly interpretable: the coordinated input direction makes no output change.

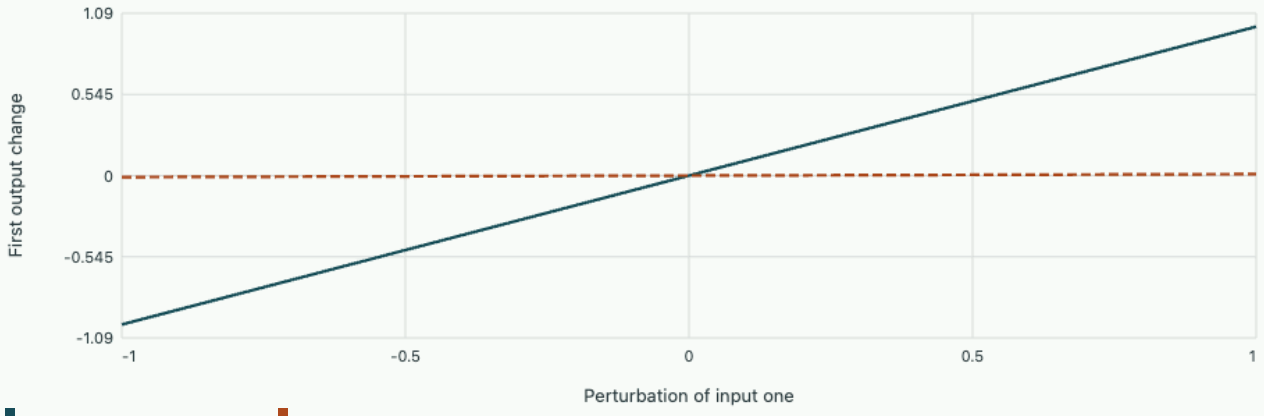
Open the live experiment

Try the calculation

Cross gain k_{12} **0.9**

Cross gain k_{21} **1.1**

Direct gain k_{22} **1**



The constrained gain is 0.0100. The diagonal relative gain is 100.00. For a unit first input the second actuator must move -1.100 units to keep its output fixed.

A change of less than the uncertainty of a single gain measurement can thus reverse the preferred feedback direction. If the second actuator reaches a limit during the constrained test, the measured slope shifts back toward the first curve because the supposedly fixed second output can no longer be maintained. Record both actuator trajectories and the residual at the constrained output before claiming to have measured a relative gain.

Tutorial · Infer the hidden actuator motion

A two-zone stage has gain matrix with rows (3, 1) and (2, 2), in consistent scaled units. Increase the first actuator by one unit while regulating the second output exactly. Find the first output and relative gain.

▼ **Write the constrained row first**

The second row requires $du_2 = -1$. The first output changes by $3 - 1 = 2$. With the second input fixed it would have changed by three, so the first relative gain is 1.5. The complementary entries are minus one half, since row and column sums must equal one. For a slowly regulating second loop the early transient can resemble the fixed-input experiment while the settled change follows the constrained one. Compare time scales before using the settled gain in a fast-loop tuning rule.

Tutorial · Quantify a near-singular inference

Suppose only the lower-right coefficient of the nearly singular example is uncertain by plus or minus 0.02. Can the sign of the diagonal relative gain be inferred from the nominal values?

▼ **Bound the determinant before taking the ratio**

At the endpoints the determinant becomes minus 0.01 and plus 0.03. At 0.99 for that coefficient the matrix is exactly singular, so the interval crosses a pole of the calculation. The sign is not identified. A covariance estimate or repeated constrained tests may narrow the uncertainty. Taking an arithmetic average of very large positive and negative sampled gains around a singularity is especially misleading; the two sides correspond to qualitatively different controller directions.

Tutorial · Test a proposed decoupling command

For the first symmetric matrix, request changes (1, minus 1) in the two temperatures. Compare the required inputs with those for (1, 1). Why is the independent-channel approximation misleading?

▼ **Invert the matrix along its two physical directions**

The equal direction has gain three and needs input changes (one third, one third). The differential direction has gain one and needs (1, minus 1). Treating both diagonal gains as two would predict changes of magnitude one half for either case. A common-mode validation therefore overstates ability to control a gradient. Test both common and differential directions at suitable amplitudes before reporting usable output ranges, especially with bounded actuators.

ETH Zürich · Control Systems II, multivariable systems and control design

ETH Zürich · MIMO directionality, singular values and relative gain

System dynamics, feedback and estimation · 3. Predicting a delayed process and testing a Smith controller

A heater is upstream of a temperature sensor. Its internal thermal response begins immediately while a material parcel reaches the sensor several seconds later. Build a feedback signal from two model outputs, derive which transfer function loses its delay, and then deliberately introduce a delay mismatch to see why a plausible prediction may fail.

Keep the measured and predicted signals distinct

Represent a stable process as a rational delay-free part $G_0(s)$ followed by transport delay e^{-Ls}. Let the controller see a prediction $y_p = y + \hat{G}_0 u - \hat{G}_0 e^{-\hat{L}s} u$. The two model terms are the currently predicted response and a stored, delayed prediction. On a perfect model, the delayed part cancels the measured response to the command, leaving an undelayed estimate in the feedback path. The measured output remains at the physical sensor downstream. Nothing in the arrangement advances an unknown disturbance in time.

With controller $u = C(r - y_p)$ and actual command-to-output plant P, direct substitution gives

$$\frac{Y}{R} = \frac{PC}{1 + C[\hat{G}_0 + P - \hat{G}_0 e^{-\hat{L}s}]},$$

$$\left.\frac{Y}{R}\right|_{\text{exact}} = e^{-Ls}\frac{G_0 C}{1 + G_0 C}.$$

The characteristic denominator is delay-free in the exact model. The reference-to-measurement signal still includes the full delay. Internal stability requires sensible models and implementation. In particular, cancelling an unstable physical mode in a transfer expression can hide a divergent internal signal. This elementary design assumes a stable plant and meaningful initial conditions for the stored model trajectory.

Calculate the entire nominal step before adjusting a tuning knob

Take $G_0 = 1/(4s + 1)$, proportional gain two and a delay of three seconds. The measured unit step response is zero until three seconds. Afterwards it is $(2/3)[1 - \exp(-3(t - 3)/4)]$. Its final offset is one third because a proportional controller is being used. It attains 63.2 percent of its own final change at 4.333 seconds after the reference step and about 95 percent at seven seconds. A plotted prediction rising before three seconds is perfectly consistent with a real sensor trace that remains flat. Labelling a model trajectory as a sensor measurement would conceal a missing or jammed sensor.

Now suppose the true transport delay is four seconds while the model still stores three. A sinusoid of angular frequency ω gives a mismatch between the two unit-magnitude delay factors equal to $2|\sin[\omega(L - \hat{L})/2]|$. At one radian per second and one second timing error this is 0.959, even though the steady mismatch vanishes. Multiply by the plant gain at that frequency before interpreting its effect. Time stamps, sample-and-hold, filtering and passage distance each contribute phase. Fitting a single dead time over a narrow frequency band cannot tell them apart.

Use the ordinary delayed-loop phase-margin study to contrast physical phase lag with the model-corrected feedback path.

Track an unknown load applied at the measured output

Let $y = Pu + d$, where a load disturbance is added after the physical process. In the exact predictor the feedback observation becomes $G_0 u + d$. With reference zero it follows that $u = -Cd/(1 + CG_0)$, so the actual output responds according to $1 - PC/(1 + CG_0)$. An abrupt output disturbance appears immediately at the sensor, but the compensating physical response arrives only after transport. Input disturbances occurring upstream have a different transfer path. Sketching the injection point avoids confusing these two bench tests.

Measure the transport delay by recording a known actuator perturbation and the two temperatures on a common clock. A feature appearing at an electronic sensor downstream may include sensor lag as well as travel. Move the sensor a known distance without changing flow and repeat. A shift proportional to added distance is evidence for transport; a constant extra lag suggests sensor or actuator dynamics. Repeat at another flow rate and revise the model if the material speed has changed.

Distinguish the two delayed model buffers in a simulation

The experiment stores the output history of the first-order lag. One buffer index represents the actual transit time and another supplies the delayed copy subtracted by the predictor. When the two delay sliders match, the difference between the physical measured curve and its ideal reference is limited to the small time-discretisation error. Change only the actual delay to test mismatch. The correction begins when one delayed signal has arrived and the other has not. Watch the predicted feedback signal separately from the physical output across that interval.

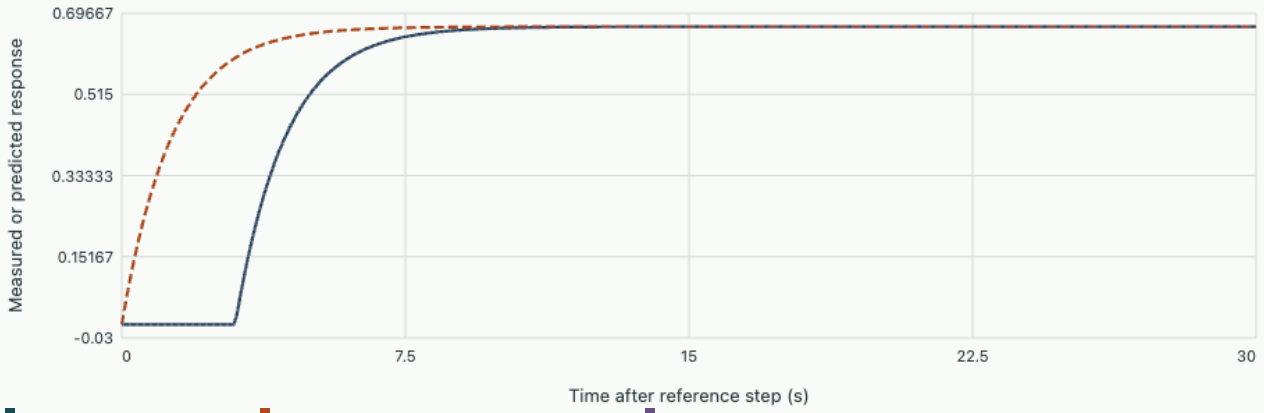
Open the live experiment

Try the calculation

Actual transport delay (s) 3

Model transport delay (s) 3

Proportional gain 2



The largest measured deviation from an exact-delay benchmark during 30 s is 0.001. The driver updates every 0.01 s. No physical response is possible before 3.0 s.

Raise proportional gain at the mismatched delay. The physical transit time stays fixed but the model-based correction becomes more sensitive to cancellation errors. A measurement which shows oscillatory excursions in this demonstration can still settle to the nominal final offset, since a pure delay has unit gain at zero frequency. Agreement at the final value cannot by itself verify the timing model.

Tutorial · Distinguish two time constants on a response plot

Two labelled traces reach 0.42 output units at 1.33 s and 4.33 s respectively after a unit step. A technician calls the second plant three times slower. Explain both traces in the nominal example and design a second check.

▼ **Align the response origins**

The 63.2 percent level of the 0.667 final value is about 0.421. An undelayed prediction reaches it at the closed-loop time constant of 1.33 s. The measured trace reaches it one transport interval later at 4.33 s. Shift the measured trace three seconds earlier and compare its entire shape, amplitude and late drift. Agreement at one level alone would also occur for many models combining a different lag and delay.

Tutorial · Quantify a frequency at which a delay error matters

For an error of 0.2 seconds, at what first positive angular frequency does the difference of unit delay factors reach magnitude one? Is that automatically a stability boundary?

▼ **Solve the phasor geometry and inspect the remaining loop**

Set $2\sin(0.1\omega) = 1$, giving $\omega = \pi/(0.6) = 5.24$ rad/s. The physical plant magnitude there is about 0.048 for the stated first-order lag. Controller gain and phase, additional modes and the nominal denominator determine the closed-loop consequence. A unit difference of raw delay phasors is an interpretable timing-error measure and does not by itself place a closed-loop pole on the imaginary axis.

Tutorial · Explain a failed transport assumption

A predictor tuned at constant flow develops oscillatory residuals when flow periodically changes. Name observations that could distinguish model-delay variation from an oscillating load.

▼ **Use independent transport and command records**

Compute the passage delay from an independent flow measurement and known pipe volume, accounting for the actual parcel history when speed varies appreciably during transit. Compare residual phase against the command and repeat with a small deliberate input perturbation at a different frequency. A response phase tracking flow supports a varying delay. An output oscillation continuing at fixed flow with no corresponding command suggests an external load or another mode. Preserve a baseline with the predictor disabled to identify plant changes independently of the feedback algorithm.

ETH Zürich · Control Systems II, dead-time compensation and the Smith predictor

System dynamics, feedback and estimation · 4. Sensitivity, bandwidth and the integral price of disturbance rejection

A regulator attenuates a slowly varying output load yet amplifies another disturbance at a higher frequency. Separate the paths by which reference, measurement noise and physical disturbance reach the recorded variable, then derive an integral identity that explains why a stable loop can redistribute sensitivity across frequency.

Solve the loop once for three independent inputs

Let the actual output be $y = Pu + d$, the measured output be $y + n$, and the command be $u = C(r - y - n)$. Here d is an additive load at the output and n is sensor error. Division by $1 + PC$ gives

$$S = \frac{1}{1 + L},$$

$$T = \frac{L}{1 + L},$$

$$L = PC,$$

$$y = Tr + Sd - Tn,$$

$$u = CS(r - d - n).$$

Thus $S + T = 1$ as a complex identity. Magnitudes generally do not add to one. Strong rejection of a slow physical load requires small sensitivity at that frequency; the measured noise can then appear almost unchanged, with opposite sign, in the actual output. A readout of the same sensor used by the controller can hide this motion. A second independent sensor can reveal it. An input load inserted before the plant would instead have transfer PS, so always identify the disturbance injection point in a reported test.

For $L = 4/[(s + 1)(s + 2)]$, both the open and closed loops are stable. The zero-frequency sensitivity is one third. At three radians per second the numerator is $-7 + 9i$ and denominator is $-3 + 9i$, giving magnitude $\sqrt{130/90} = 1.202$. A 0.10-unit output disturbance at that frequency produces a 0.120-unit actual output oscillation. A time plot of reference following alone would miss this amplification. The sensitivity peak also measures proximity of the Nyquist curve to minus one, since $1/|S| = |1 + L|$.

Derive the stable rational case from its pole and zero distances

For a closed-loop stable rational sensitivity tending to one at high frequency, write its numerator and denominator as monic products of roots. For a real negative root at minus a, a factor contributes $\frac{1}{2}\log(\omega^2 + a^2)$. The difference for two such factors integrates by differentiating under the integral in parameter a. The derivative $a/(\omega^2 + a^2)$ has integral $\pi/2$ over the positive frequency axis. Integration back gives $\pi(a - b)/2$. For complex conjugate roots the same result follows by pairing both signs of frequency and shifting the full-axis integration variable by the imaginary part. The area contribution depends on the absolute real part.

The sensitivity zeros are the open-loop poles and its poles are the closed-loop poles in a minimal description. For relative degree at least two the coefficient of the next-to-highest power in the open- and closed-loop characteristic polynomials is unchanged; their signed root sums therefore coincide. If every open-loop pole lies in the left half-plane, the absolute-real-part sums also coincide and the integral vanishes. An open-loop pole in the right half-plane reverses its absolute-real-part contribution and leaves the general result

$$\int_0^\infty \log|S(i\omega)|\,d\omega = \pi \sum_{p\in\mathrm{RHP}} \operatorname{Re} p.$$

This statement assumes internal closed-loop stability, suitable high-frequency roll-off and no imaginary-axis poles in the elementary derivation. It is an area on a linear angular-frequency axis with natural logarithm. Estimating it as signed visual area on a conventional log-frequency Bode plot weights the frequencies differently. A stable low-pass system can reduce sensitivity over a wide low-frequency interval by allowing higher sensitivity elsewhere. The identity does not fix an individual peak without bandwidth or shape constraints.

Connect the complementary sensitivity to model error

For multiplicative plant uncertainty $P_{\text{actual}} = P(1 + W\Delta)$, factor the return difference as $(1 + L)(1 + TW\Delta)$. If nominal interconnection is internally stable, the uncertainty and weighting are stable, and $\|TW\|_\infty < 1$, the small-gain theorem ensures stability for all such perturbations with $\|\Delta\|_\infty \le 1$. A failure of this sufficient check invites closer analysis of the particular physical uncertainty. A poorly known resonance above bandwidth can matter even after a nominal step trace has settled smoothly. State the frequency range on which input-output identification actually constrains the model.

Change a lead network and compare crossover with measured sensitivity near a resonance.

Locate the sensitivity peak and test vector addition

Sweep the gain numerator at fixed pole locations. At zero gain the load sees essentially no correction. Increasing gain initially improves low-frequency load rejection but can create a distinct sensitivity maximum between the pole scales and the high-frequency range. Find the frequency where the sensitivity curve crosses one. At that point the distance of the loop phasor from minus one equals one. The complementary-sensitivity magnitude there need not vanish; the sum identity concerns complex values with their phases retained.

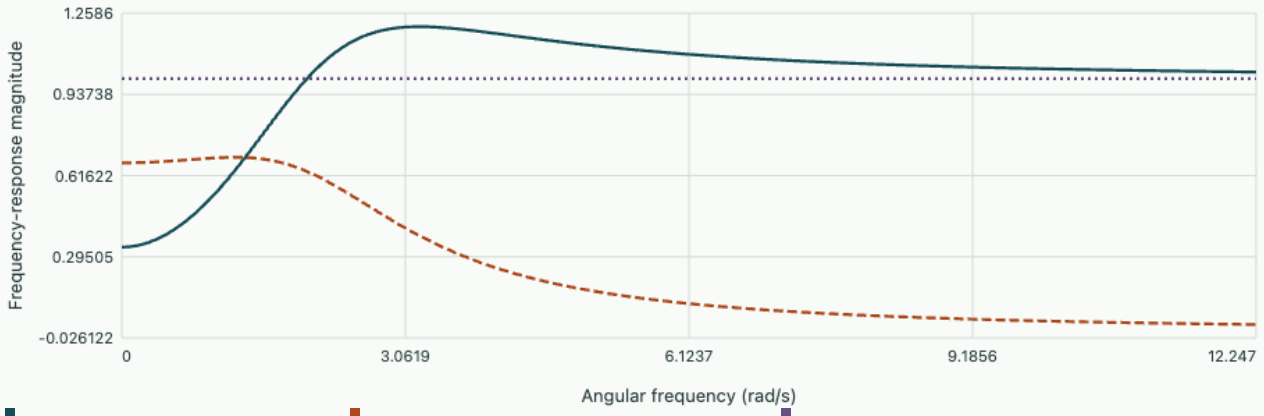
Open the live experiment

Try the calculation

Loop numerator K **4**

First pole frequency (rad/s) **1**

Second pole frequency (rad/s) **2**



Static sensitivity is 0.333. The sampled peak is 1.206 near 3.20 rad/s. At 3 rad/s |S| = 1.202.

To check the curves on a real apparatus, inject small reference tones to estimate complementary sensitivity and separate, calibrated output loads to estimate sensitivity. Record coherent amplitude and phase at matching frequencies with the same time origin. Sum the two complex estimates and compare against unity. Delay in the reference measurement chain alone can spoil that check, even if both measured magnitudes seem plausible. The test is most informative where both responses have appreciable amplitude.

Tutorial · Test an apparent counterexample

For $L = K/(s + a)$ with positive K, a, sensitivity is below one at all finite frequencies. Evaluate the logarithmic area and locate the failed assumption.

▼ **Integrate the log-distance difference directly**

The sensitivity equals $(s + a)/(s + a + K)$. The real-root identity gives

$$\pi[a - (a + K)]/2 = -\pi K/2.$$

This loop has relative degree one. Its numerator and denominator root sums differ by K, invalidating the high-frequency cancellation used for the zero-area result. Adding a finite fast actuator pole restores relative degree two and can move the compensating sensitivity region far beyond the frequencies first

plotted. Check both the analytical degree and measurement bandwidth before diagnosing a contradiction.

Tutorial · Use two independent sensors

At a low frequency the loop has real positive gain 19. A calibrated sinusoid of amplitude 0.20 mm is electronically added to the feedback sensor with no physical load. Predict the actual motion and its indicated motion at that sensor.

▼ **Follow the noise to two separate outputs**

Sensitivity is 0.05 and complementary sensitivity 0.95. The actual stage moves with amplitude 0.19 mm opposite the injected signal. Its indicated position at the corrupted sensor has amplitude 0.01 mm in phase with the injection. An independent reference encoder sees the much larger actual motion. Increasing gain further would make the feedback readout look quieter while causing more physical motion in this particular test.

Tutorial · Bound uncertainty at a measured spectral line

At one frequency the nominal measured complementary sensitivity is 1.25 and the magnitude bound on multiplicative uncertainty is 0.60. At a second frequency the values are 1.05 and 1.10. Interpret both products.

▼ **Apply the pointwise ingredient of the small-gain check**

The products are 0.75 and 1.155. The first frequency meets the sufficient magnitude bound there. The second exceeds it, so a claim based on that global criterion is unavailable. A tighter phase-aware model or structured uncertainty can refine the conclusion. Checking two isolated frequencies cannot establish a supremum bound, since resonant features may occur between samples. Estimate frequency resolution from the narrowest expected damping bandwidth and verify both sides of each measured peak.

Åström and Murray · Feedback Systems, chapters on sensitivity functions and Bode's integral formula

System dynamics, feedback and estimation · 5. Which sampling times reveal the velocity of an oscillator?

A sensor measures displacement while a controller needs velocity as well. Use the exact free vibration between two readings to recover the missing state, then quantify how a poorly chosen sample interval amplifies measurement noise.

Observe a rotation through one coordinate

For an undamped, unforced oscillator $\ddot{x} + \omega^2 x = 0$, use the state $z = [x, \dot{x}/\omega]^T$. Both entries then have units of length. A delay Δt advances phase by $\theta = \omega \Delta t$. Integrating the equation exactly gives

$$x(t + \Delta t) = x(t)\cos\theta + [\dot{x}(t)/\omega]\sin\theta.$$

The velocity equation follows by differentiating with respect to Δt. Together they rotate the state by a matrix with first row $(\cos\theta, \sin\theta)$ and second row $(-\sin\theta, \cos\theta)$.

The position sensor has observation row $(1, 0)$. At consecutive samples its noiseless readings obey

$$\begin{bmatrix} y_0 \\ y_1 \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 0 \\ \cos\theta & \sin\theta \end{bmatrix}}_{O} \begin{bmatrix} x_0 \\ v_0/\omega \end{bmatrix}.$$

The determinant is $\sin\theta$. At integer half-period intervals the sensor measures the same direction in the state plane twice, allowing opposite velocities to produce exactly the same pair of readings. Away from those intervals,

$$v_0 = \omega(y_1 - y_0\cos\theta)/\sin\theta.$$

The determinant diagnoses whether an exact reconstruction exists. The singular values describe whether that reconstruction will tolerate small errors. Because OO^T has diagonal entries one and off-diagonal entries $\cos\theta$, its eigenvalues are $1 \pm |\cos\theta|$. Their square roots are the singular values of O in the chosen state scaling.

Propagate the uncertainty in two readings

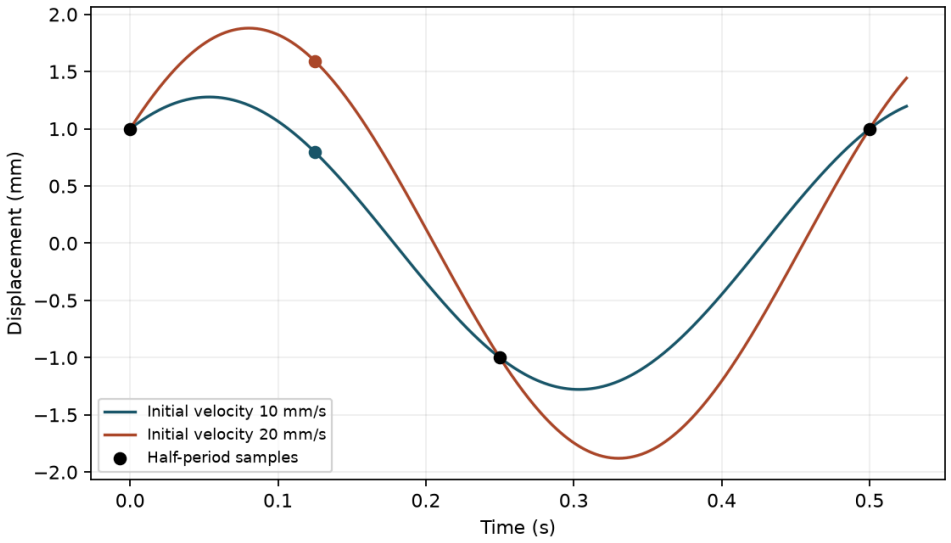
Suppose each recorded position has independent zero-mean noise of standard deviation σ_x. A linear combination $y_1 - y_0\cos\theta$ has variance $\sigma_x^2(1 + \cos^2\theta)$. Multiply by the square of the velocity gain to obtain

$$\sigma_v = \omega\sigma_x\sqrt{1 + \cos^2\theta}/|\sin\theta|.$$

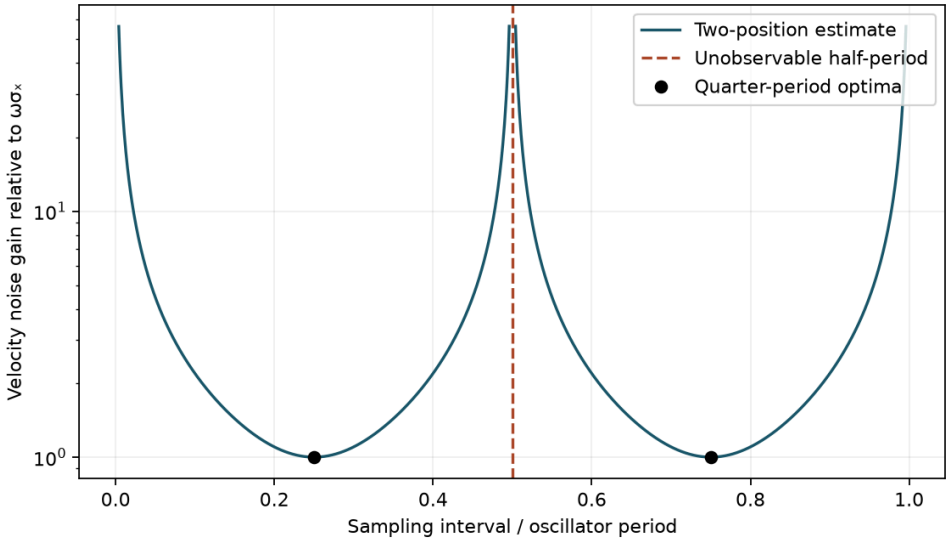
An exact knowledge of the sample times and natural frequency is assumed in this expression.

A 2 Hz oscillator has period 0.5 s and angular frequency 4π rad/s. With a 0.125 s interval the second reading directly measures v_0/ω. For 0.10 mm position noise the recovered velocity has standard

deviation 1.257 mm/s. Increasing the delay to 0.240 s moves the second observation close to the direction opposite the first. The same position uncertainty now produces 14.12 mm/s of velocity uncertainty, although the mathematical reconstruction still exists. The notebook varies the sampling interval continuously and checks the formula against noise propagation through the inverse matrix.



Compare the equal half-period readings with the selected off-grid observation. Continuous motion between sample points can carry information that the readings discard.



The smallest uncertainty occurs at the quarter-period points where the second position reading depends only on the initial velocity.

Problem 1 · Take more readings at exactly one half-period

A colleague proposes collecting 100 equally spaced positions to recover velocity at $\Delta t = T/2$. Does the larger data set recover the missing state for this exact model?

▼ Work through the solution

After k steps the observation row is $(\cos k\pi, \sin k\pi) = ((-1)^k, 0)$. Every reading measures only the initial position, up to its alternating sign. Averaging can reduce the noise in x_0, while the initial velocity

remains arbitrary. One extra measured position at a known quarter-period after the first would add row $(0, 1)$ and restore the missing independent equation.

Problem 2 · Relate observability to a finite difference

For $\omega \Delta t \ll 1$, derive the leading noise in the reconstructed velocity. Explain why decreasing the interval can increase the velocity error here.

▼ **Work through the solution**

Use $\sin\theta \simeq \theta = \omega\Delta t$ and $\cos\theta \simeq 1$. The estimate becomes $(y_1 - y_0)/\Delta t$ to leading order and $\sigma_v \simeq \sqrt{2}\sigma_x/\Delta t$. The actual displacement accumulated over a short interval shrinks, while the two independent noise samples have a difference of fixed standard deviation. This result applies to the two-reading estimate; fitting many accurately timed readings over a fixed observation duration uses different information.

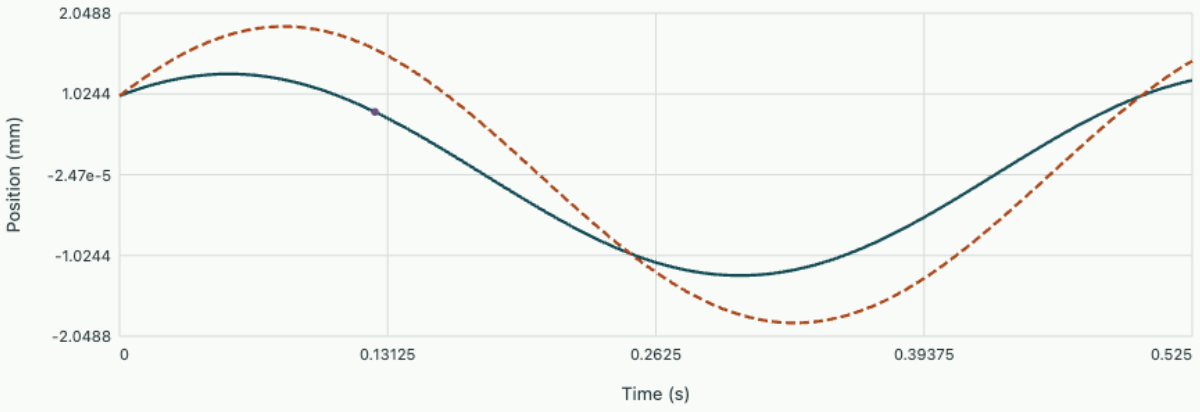
Interactive calculation · Oscillator sampling

Change the physical and numerical inputs of the notebook model. Compare the updated curves and balance checks with the recorded notebook figures. Each plot uses the same equations and assumptions documented in that notebook; its readout names any analytical or reduced-model reference.

Open this live calculation

Try the calculation

Natural frequency (Hz)	2
Initial velocity (mm/s)	10
Sampling / period	0.25



For independent0.10 mm position noise the inferred velocity standard deviation is 1.257 mm/s.

Download this example’s editable notebook

Interactive calculation · Sampling noise

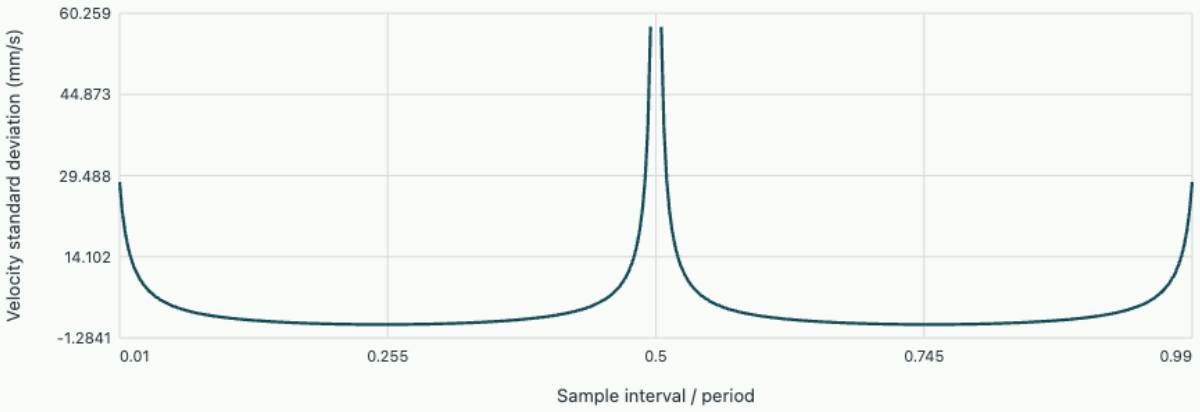
Change the physical and numerical inputs of the notebook model. Compare the updated curves and balance checks with the recorded notebook figures. Each plot uses the same equations and assumptions documented in that notebook; its readout names any analytical or reduced-model reference.

Open this live calculation

Try the calculation

Natural frequency (Hz) **2**

Position noise standard deviation (mm) **0.1**



Two-position reconstruction

Quarter-period minimum standard deviation 1.257 mm/s. The divergence near half-period and full-period sampling is a loss of observability, not floating-point noise.

Download this example's editable notebook

Open editable notebook Download .ipynb

System dynamics, feedback and estimation · 6. When does a stable operating point disappear?

A shaft can hold a static torque against gravity up to a finite limit. Trace the equilibrium branches and distinguish local stability, loss of an equilibrium and the possibility of oscillation.

Construct the operating diagram from a torque balance

A pendulum driven by a constant torque u obeys $J\ddot{\theta} + d\dot{\theta} + k\sin\theta = u$, where $k = mg\ell$, $d > 0$ and $J > 0$. With $v = \dot{\theta}$ the state arrows are $(v, (u - dv - k\sin\theta)/J)$. The first nullcline is $v = 0$. The second is $v = (u - k\sin\theta)/d$. Plotting both gives the local arrow directions before any time integration. Their intersections are precisely static operating points.

For $|u| < k$ there are two intersections per revolution: $\theta_s = \arcsin(u/k)$ and $\theta_u = \pi - \theta_s$, modulo 2π. Linearize about an intersection to obtain

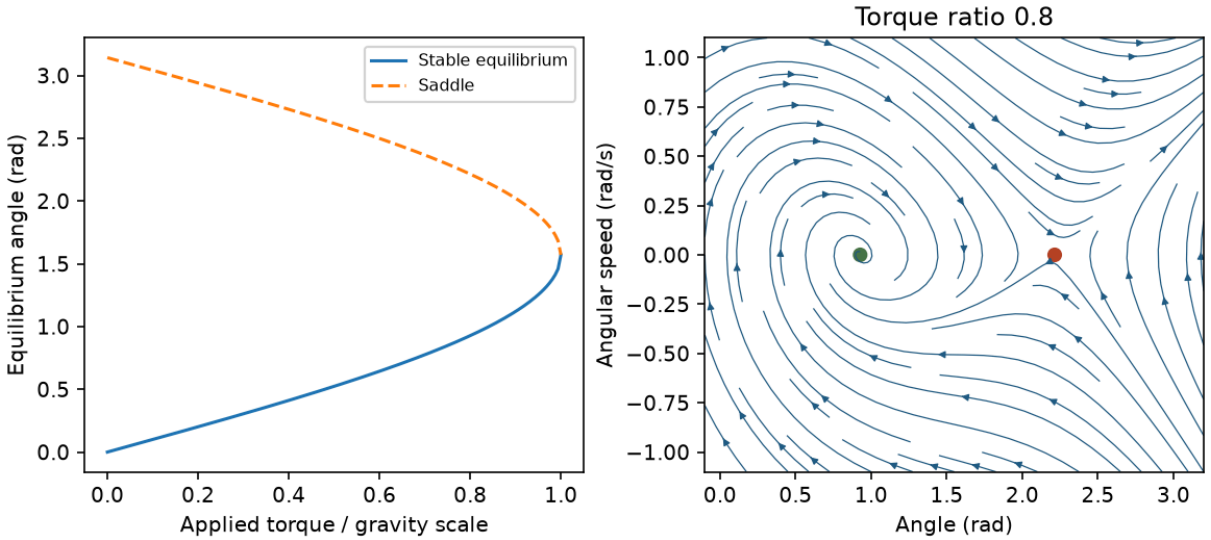
$$J\lambda^2 + d\lambda + k\cos\theta_e = 0.$$

Positive incremental restoring stiffness gives stable poles for $d > 0$. Negative stiffness gives roots of opposite signs and therefore a saddle. At $u = k$ the branches meet and one pole is zero. For $u > k$ there is no static balance. A gain tuned from a linearization on the former branch cannot hold a nonexistent open-loop equilibrium at the same constant torque.

Recover the saddle-node scaling without fitting data

Set $\theta = \pi/2 + x$ and $u = k + \mu$. Since $\sin(\pi/2 + x) = 1 - x^2/2 + O(x^4)$, the slow balance near the merging points is $d\dot{x} \simeq \mu + kx^2/2$. The omitted acceleration changes the fast transient; at the merger one eigenvalue remains $-d/J$ while the other vanishes. For $\mu < 0$ the displacement from the merger scales as $|x| \simeq \sqrt{-2\mu/k}$. The restoring torque gradient tends to zero and perturbations take longer to decay. This is a fold or saddle-node bifurcation.

Use $J = 1$ kg m², $d = 0.4$ N m s/rad, $k = 1$ N m and $u = 0.8$ N m. The equilibria are 0.9273 and 2.2143 rad. Their stiffnesses are +0.6 and −0.6 N m/rad. The stable poles are $-0.2 \pm 0.748i$ s⁻¹. The saddle poles are $+0.6$ and -1.0 s⁻¹. A simulated trajectory that sits exactly at the saddle can stay there indefinitely in exact arithmetic. Perturb it along either eigenvector to identify its phase-plane structure.



The stable and saddle branches meet at the load limit. The phase portrait at 80% load shows trajectories attracted to the lower-angle equilibrium and deflected from the saddle.

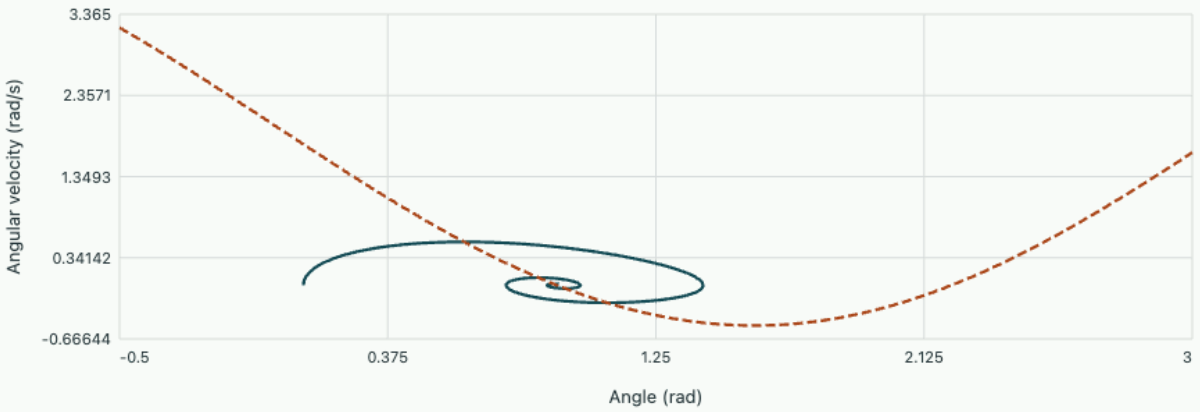
Interactive calculation · Move through the loss of a static operating point

Increase the applied torque to and beyond the gravity limit. Change the starting angle to approach either side of the saddle. The browser integrates the actual pendulum equation with RK4, using inertia1 kg m² and gravity-torque scale1 N m.

Open this live calculation

Try the calculation

Applied / maximum static torque	**0.8**
Damping (N m s/rad)	**0.4**
Initial angle (rad)	**0.1**



Trajectory from rest Acceleration nullcline

Stable branch angle 0.927 rad; saddle branch 2.214 rad. At unity load the branches meet. Inertia and gravity torque scale are both one.

Compare a fold with the onset of a periodic motion

A local stability change need not destroy the equilibrium. In the illustrative planar normal form $\dot{x} = \mu x - y - (x^2 + y^2)x$, $\dot{y} = x + \mu y - (x^2 + y^2)y$, put $r^2 = x^2 + y^2$. Differentiation gives $\dot{r} = \mu r - r^3$ and $\dot{\varphi} = 1$. The origin persists on both sides. For $\mu > 0$ it loses stability and a stable periodic orbit of radius $\sqrt{\mu}$ appears. This supercritical Hopf example has a fundamentally different phase portrait from the torque fold.

What a sweep should record

A bifurcation diagram needs an observable appropriate to each branch. Record equilibrium angles after transients for static branches and maxima and minima for cycles. Direct forward simulation usually misses unstable branches. Solving the equilibrium equations and inspecting their Jacobian recovers them.

Exercises on stability evidence

At $u/k = 0.99$, use the fold expansion to estimate the angular separation of the two equilibria.

▼ **Worked solution**

The near-merger displacement is $\sqrt{2(0.01)} = 0.1414$ rad on either side. The separation is 0.2828 rad, close to the exact $2\arccos(0.99) = 0.2831$ rad.

At $\mu = 0.04$ in the planar normal form, what does a radial perturbation to the periodic orbit do?

▼ **Worked solution**

The orbit radius is 0.2. The radial derivative $\mu - 3r^2 = -0.08$ gives local decay with time constant 12.5 model time units. A phase shift along the orbit persists.

Further reading · Åström and Murray, Feedback Systems, 2nd ed., Chapter 5.

System dynamics, feedback and estimation · 7. Can any controller meet this bandwidth and peak-gain specification?

Delay uses up phase at the crossover even when it changes no gain. An unstable pole near a right-half-plane zero imposes a separate peak-gain limit. Derive each restriction before tuning a controller.

Spend a phase budget at the gain crossover

In a scalar negative-feedback loop let $L = PC$, $S = 1/(1+L)$ and $T = L/(1+L)$. At $|L(i\omega_c)| = 1$, a required margin ϕ_m imposes $\arg L(i\omega_c) \geq -\pi + \phi_m$ on the continuously unwrapped phase at the crossing relevant to stability. Factor a right-half-plane zero at $z > 0$ into a minimum-phase zero and the unit-gain all-pass factor $(z-s)/(z+s)$. Its phase is $-2\arctan(\omega/z)$. Pure delay contributes $-\omega\tau$. Magnitude data alone cannot reveal these phase costs.

For a nominal minimum-phase design contributing -90° at crossover, the available phase budget gives the following exact condition under that assumption. An attainable gain crossover with multiple crossings still requires the full Nyquist stability test.

$$2\arctan(\omega_c/z) + \omega_c\tau \;\leq\; \pi/2 - \phi_m.$$

If $z = 40$ rad/s, $\tau = 0$ and $\phi_m = 45^\circ$, then $\omega_c \leq 40\tan(\pi/8) = 16.57$ rad/s. A proposed 25 rad/s crossover would leave only $90^\circ - 2\arctan(25/40) = 25.99^\circ$. A changed controller may buy lead with noise amplification or a steeper high-frequency roll-off. Reevaluate the complete minimum-phase contribution rather than treating this particular 16.57 as universal.

Use the maximum-modulus principle for a controller-independent test

Now suppose the plant has a simple real unstable pole $p > 0$ and a right-half-plane zero $z > 0$ with $p \neq z$. Assume rational proper internally stabilizing feedback without unstable cancellations. S and T are bounded and analytic in the open right half plane. Internal stability of PS requires $S(p) = 0$, hence $T(p) = 1$. At the plant zero $T(z) = 0$. Divide T by the stable all-pass factor $B_z(s) = (s-z)/(s+z)$. The apparent singularity at z is removable. On the imaginary axis $|B_z| = 1$, so $F = T/B_z$ and T have identical boundary magnitudes.

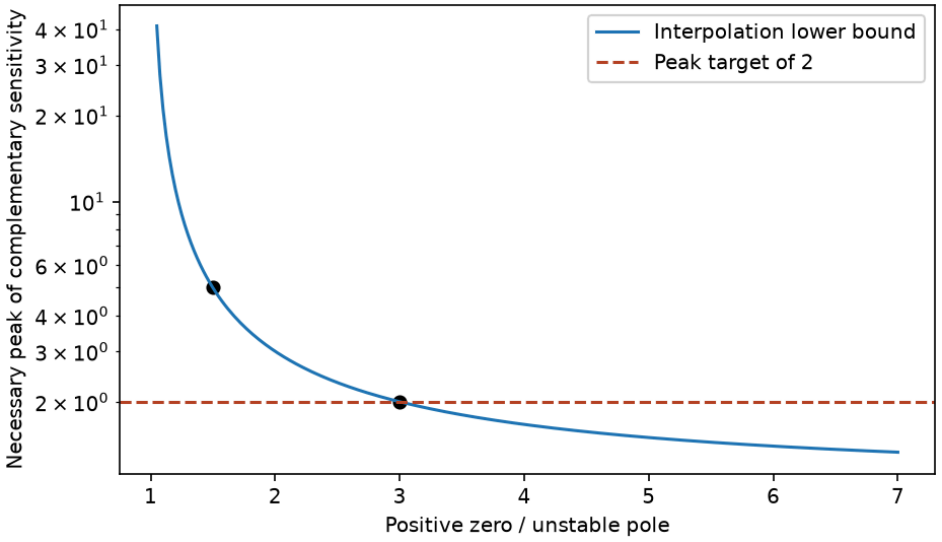
The maximum-modulus principle bounds the interior value $|F(p)|$ by the largest boundary magnitude (including the limit at infinity). Substituting the interpolation values gives

$$\|T\|_\infty \geq |F(p)| = \left|\frac{p+z}{p-z}\right|.$$

For $p = 2$ and $z = 3$ s⁻¹ the peak is at least 5, approximately 14.0 dB. A claimed peak below 2 is infeasible for every controller obeying the stated assumptions, even if the measured range missed the

peak. When the pole and zero approach each other the bound diverges. Trying to cancel them inside the controller sacrifices internal stability.

> **Why the sensitivity integral and this test say different things**
>
> The Bode integral constrains the area of log sensitivity across frequency. The maximum-modulus argument uses exact internal interpolation at two specific plant singularities to bound a peak of complementary sensitivity. Which result is sharper depends on the requirement and measured plant.



The necessary peak diverges as the unstable pole and positive zero approach each other. A zero-to-pole ratio of three is the first compatible with this necessary bound and a peak target of two.

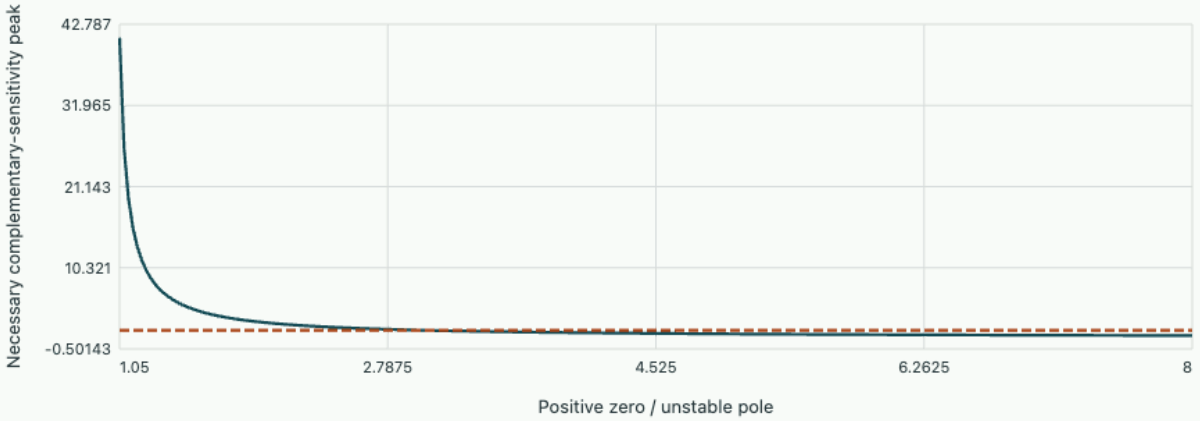
Interactive calculation · Move a positive zero toward an unstable pole

Compare the necessary peak-gain bound with the specified target. The delay affects the separate45° phase budget, not the interpolation inequality.

Open this live calculation

Try the calculation

Unstable pole p (1/s)	1
Allowed peak gain	2
Transport delay (s)	0.05



The positive zero must satisfy z/p ≥ 3.000. At z=3p, a 45° phase budget gives crossover at most 1.143 rad/s. Both conditions are necessary, not a controller construction.

Exercises on impossible requirements

A measured minimum-phase loop has −110° phase at its proposed crossover of 30 rad/s. How much additional pure delay permits a 45° margin?

▼ **Worked solution**

The spare phase is $180 - 110 - 45 = 25^\circ = 0.4363$ rad. At the stated unchanged crossover this permits at most $0.4363/30 = 14.5$ ms.

How far must a positive zero $z > p$ lie above an unstable pole $p = 2$ to avoid ruling out a peak target of 2 by the interpolation bound alone?

▼ **Worked solution**

Require $(z+2)/(z-2) \leq 2$, giving $z \geq 6$ s⁻¹. This removes one necessary-condition violation. It does not construct a controller satisfying every constraint.

Further reading · Åström and Murray, Feedback Systems, 2nd ed., Chapters 13–14, especially sections 14.3–14.4.

System dynamics, feedback and estimation · 8. Designing current, speed and position loops for a drive

An encoder measures shaft motion, while the power stage acts on voltage. Decide which disturbances each measurement can reject and how fast each loop can move before the unseen inner dynamics erode the position margin.

Place a fast loop where the disturbance enters

Use the small-signal motor equations $L\dot{i} + Ri = v - k_e\omega$, $J\dot{\omega} + B\omega = k_t i - \tau_L$, and $\dot{\theta} = \omega$. Measure current at the converter and motion at the shaft. A speed loop alone leaves rapid electrical disturbances to propagate before they appear in shaft motion. Close a fast local current loop and send it a current reference from the mechanical loop. Treat back EMF as a disturbance or add measured-speed feedforward.

Choose $C_i(s) = k_{pi} + k_{ii}/s$ with $k_{ii}/k_{pi} = R/L$. For the nominal $1/(Ls + R)$ current plant the open loop reduces to $k_{pi}/(Ls)$ and the closed transfer is $\omega_i/(s + \omega_i)$, with $k_{pi} = L\omega_i$ and $k_{ii} = R\omega_i$. At $L = 0.010$ H, $R = 1$ ohm and $\omega_i = 1000$ rad/s the gains are 10 V/A and 1000 V/(A s). The approximation presupposes adequate measurement, switching and computation bandwidth.

Choose the middle loop against an inner-loop error budget

If current follows its command over the speed-loop band, the torque-to-speed plant is $k_t/(Js + B)$. A speed PI controller with $k_{p\omega} = J\omega_v/k_t$ and $k_{i\omega} = B\omega_v/k_t$ gives an approximate closed speed response $\omega_v/(s + \omega_v)$. With $J = 0.020$, $B = 0.020$, $k_t = 0.40$ in SI and $\omega_v = 100$ rad/s, the gains are 5 in the corresponding current-per-speed and integral units. The neglected current lag contributes $\arctan(100/1000) = 5.71^\circ$ at that speed-loop crossover. Reserve this phase when validating the full model. A torque disturbance enters the mechanical balance directly and is primarily rejected by this middle loop.

Convert a position requirement into a realizable crossover

Use position proportional feedback $\omega_{ref} = K(\theta_{ref} - \theta)$, the achieved speed model $\omega_v/(s + \omega_v)$ and total outer-path delay τ. In this reduced model the exact position-loop frequency response is $L(s) = Ke^{-s\tau}/[s(1 + s/\omega_v)]$. Choosing K for a desired amplitude crossover separates magnitude tuning from the residual phase budget:

$$K = \omega_c\sqrt{1 + (\omega_c/\omega_v)^2},$$

$$\phi_m = \pi/2 - \arctan(\omega_c/\omega_v) - \omega_c\tau.$$

At $\tau = 6$ ms, $\omega_v = 100$ rad/s and required margin 45°, solving the monotone equation gives $\omega_c \le 51.56$ rad/s. The 50 rad/s design has only 46.25° before any additional sensor or drive lag. At 20 rad/s the nominal margin is 71.81°. The notebook separately evaluates the delay-free closed-loop poles

with OneSimulation. Their decay rate and the delayed loop's phase margin measure different properties.

Interactive experiment · Spend the phase budget with bandwidth and delay sliders

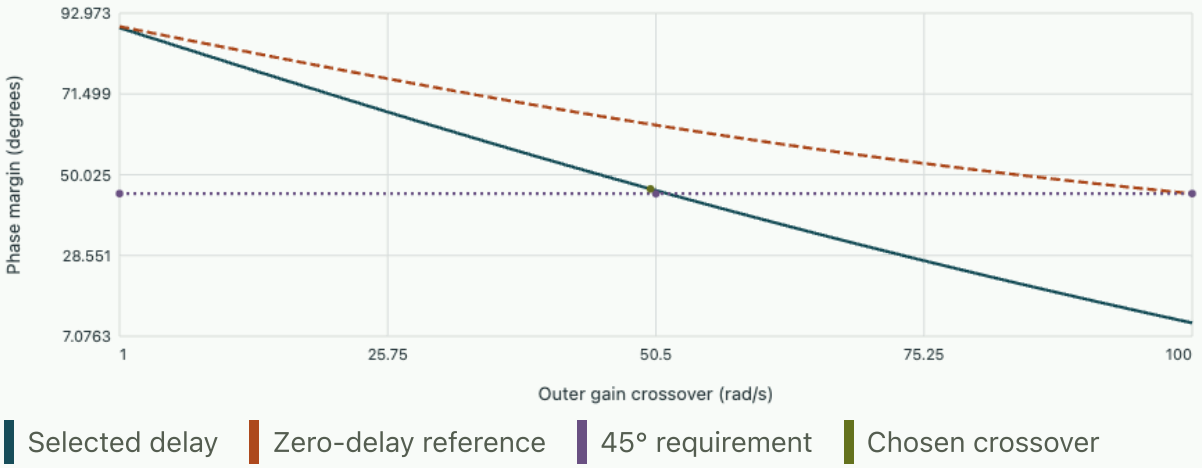
Choose the achieved inner-loop bandwidth, the outer-path delay and the intended outer crossover. The browser adjusts the proportional gain to place the amplitude crossover exactly at that frequency, then evaluates the remaining phase margin. The 45° line is a design target; it is not the boundary of stability itself.

Begin at an inner bandwidth of 100 rad/s, delay 6 ms and crossover 50 rad/s. Double the delay and observe the lost 17.19°. Then lower the requested crossover until the target is recovered. The zero-delay pole margin in the readout is computed from the quadratic closed-loop characteristic polynomial. Compare it with the OneSimulation stability diagnostic in the notebook; it does not certify the delayed system.

Open this experiment online

Try the calculation

Achieved inner bandwidth (rad/s)	**100**
Outer-path delay (ms)	**6**
Chosen outer crossover (rad/s)	**50**



Chosen gain K = 55.902 /s gives 46.25° phase margin. The 45° boundary is 51.56 rad/s. The delay-free closed-loop pole margin is 50.000 /s; it is a separate diagnostic from the delayed phase margin.

Continue in the OneSimulation notebook · Download its Python model

OneSimulation experiment · An achieved inner speed bandwidth and transport delay bound the outer-loop crossover

Run the notebook with OneSimulation 0.2.0, NumPy, SciPy and Matplotlib. Change the input data and compare your run with the recorded result.

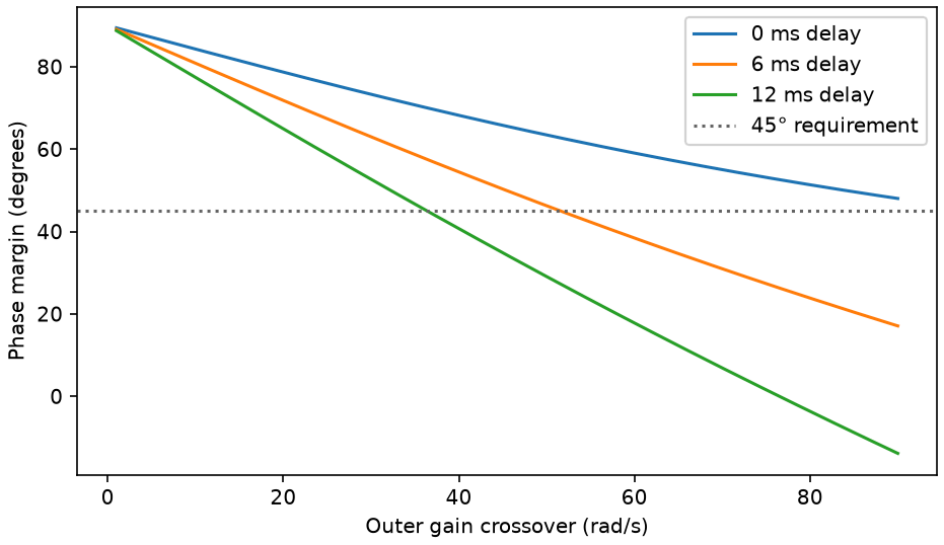
```
"""Evaluate inner-lag and transport-delay limits on an outer position loop."""
import numpy as np
import matplotlib.pyplot as plt
from onesim.physics.systems.control import stability_margin

inner_pole = 100.0  # rad/s, achieved bandwidth of the closed speed loop
def outer_gain(crossover):
    return crossover*np.sqrt(1+(crossover/inner_pole)**2)
def phase_margin(crossover, delay):
    return 90-np.rad2deg(np.arctan(crossover/inner_pole)+crossover*delay)

frequencies = np.linspace(1,90,250)
fig, ax = plt.subplots(figsize=(6.4,3.7), layout="constrained")
for delay in [0.0,0.006,0.012]:
    ax.plot(frequencies, phase_margin(frequencies,delay),label=f"{1000*delay:.0f} ms delay")
for crossover in [20.0, 50.0]:
    gain = outer_gain(crossover)
    poles = np.roots([1.0,inner_pole,inner_pole*gain])
    print(f"crossover {crossover:.0f} rad/s: delay-free pole margin = "
          f"{stability_margin(poles,discrete=False):.3f} /s, "
          f"6 ms phase margin = {phase_margin(crossover,0.006):.2f} deg")
ax.axhline(45,color="0.4",ls=":",label="45° requirement")
ax.set(xlabel="Outer gain crossover (rad/s)",ylabel="Phase margin (degrees)")
ax.legend()
plt.show()
```

Reference run

```
crossover 20 rad/s: delay-free pole margin = 28.543 /s, 6 ms phase margin = 71.81 deg
crossover 50 rad/s: delay-free pole margin = 50.000 /s, 6 ms phase margin = 46.25 deg
```



An achieved inner speed bandwidth and transport delay bound the outer-loop crossover

Download editable notebook · Download Python script

Inspect failure signatures from the inside out

Measure the current-reference response with the mechanical loop open in the model. Then close speed feedback and compare measured torque-step recovery with the linear prediction. Finally apply a position command and log reference and achieved current, shaft speed and position together. Under voltage clipping the actual current cannot track its command at speed. A speed integrator that keeps charging can dominate the later position excursion. Use actuator feedback for anti-windup and check the actual limits at the operating speed.

Which loop receives the position sample delay?

A rotor-position timestamp used to derive speed can delay both mechanical loops. A transport delay introduced only between the supervisory controller and the drive delays the outer loop. Draw the measurement paths before inserting a single lumped delay or counting it twice.

Exercises on a changed architecture

Keep crossover at 50 rad/s while total delay rises from 6 to 12 ms. What does this do to the reduced-model margin?

▼ **Worked solution**

The gain is unchanged by pure delay. The added lag is $50(0.006) = 0.30$ rad, or 17.19°. Margin falls from 46.25° to 29.06°, below the target.

At high speed the current sensor follows its local model but digital speed estimates have 4 ms of additional filtering delay. Where should the new uncertainty enter?

▼ **Worked solution**

Place it in the speed measurement path. Recompute the middle-loop closed dynamics before using them as the position plant. Adding it only to the outer-loop delay would miss loss of speed-loop robustness.

Further reading · Åström and Murray, Feedback Systems, 2nd ed., Chapters 10–11 and 15.