Vibration and identification in machinery

Connect modes, moving components and measured frequency responses. Investigate how forcing, damping and the observation point change an inferred dynamic model.

Contents

1. **Critical speed, synchronous whirl and two different rotor centres**
   A stationary probe observes motion of a rotating shaft surface, while a balance calculation follows an eccentric mass centre. Derive both trajectories from one flexible-rotor model. Their difference explains the apparent recentring of a rotor at high speed and why a single vibration amplitude cannot identify unbalance and damping.
2. **Parametric resonance under periodic stiffness modulation**
   A carriage supported by a spring whose stiffness changes periodically can gain oscillation energy with no additive force at its natural frequency. Compare this mechanism with an ordinary forced resonance, estimate the first instability band and learn how to measure growth before nonlinear saturation invalidates the prediction.
3. **Random base motion, spectral density and rms response**
   A support vibration is described by a power spectral density rather than a single sine-wave amplitude. Derive the relative-motion filter for a mounted instrument, calculate mean-square motion from a measured spectrum and distinguish a stationary rms prediction from a peak displacement guarantee.
4. **Why a local damper couples otherwise separate modes**
   Undamped mode shapes diagonalise a symmetric mass and stiffness model. Adding a damper at one attachment does not generally leave modal equations separate. Derive the resulting coupling, test its magnitude and learn when a diagonal modal-damping approximation can hide an observable phase error.
5. **Which damping law produced the free-decay trace?**
   A carriage slides on a guide and a spring pulls it back to its resting position. Two dampers can be adjusted to give the same first peak reduction. Predict the following turning points and decide which measurement separates viscous loss from dry sliding friction.
6. **From beam curvature to a mass matrix and a frequency measurement**
   A coarse cantilever model often predicts its first resonance accurately while putting the third far too high. Derive both matrices from the same displacement field and inspect how each mode converges.

7. Which approximation shifted the resonance?

A simulation of a pulsed cantilever contains three separate approximations. Measure the effect of spatial interpolation, retaining finitely many modes and representing a changing force on a time grid.

Vibration and identification in machinery · 1. Critical speed, synchronous whirl and two different rotor centres

A stationary probe observes motion of a rotating shaft surface, while a balance calculation follows an eccentric mass centre. Derive both trajectories from one flexible-rotor model. Their difference explains the apparent recentring of a rotor at high speed and why a single vibration amplitude cannot identify unbalance and damping.

Choose the observed coordinate before writing a force

A symmetric disk of mass m sits at the centre of a massless flexible shaft with equal lateral stiffness k in two directions and viscous damping c_d acting on its geometric centre. Its mass centre lies a distance e from that point. Assume constant spin Ω, small lateral deflections, simple supports and negligible gyroscopic tilt. Write the complex displacement of the shaft centre as $z = x + iy$. The mass centre has displacement $z + ee^{i\Omega t}$. Its acceleration enters Newton's law while bearing restoring and damping forces depend on the shaft-centre motion:

$$m\ddot{z} + c_d\dot{z} + kz = me\Omega^2 e^{i\Omega t}.$$

Gravity merely shifts the static position for a horizontal isotropic linear shaft and can be subtracted before fitting the synchronous response. Set $\omega_n = \sqrt{k/m}$, $\zeta = c_d/(2m\omega_n)$ and $r = \Omega/\omega_n$. The steady response divided by eccentricity is

$$\frac{Z}{e} = \frac{r^2}{1 - r^2 + 2i\zeta r},$$

$$\frac{Z_{\mathrm{mass}}}{e} = 1 + \frac{Z}{e} = \frac{1 + 2i\zeta r}{1 - r^2 + 2i\zeta r}.$$

At small r the shaft hardly moves while the mass centre describes its geometric eccentricity. Near resonance the shaft amplitude rises and lags the unbalance reference by about a quarter turn. At large r shaft motion approaches amplitude e in opposite phase to the eccentricity, leaving the mass centre comparatively close to the bearing line. A shaft probe does not then read zero merely because the mass centre recentres.

Measure amplitude and phase with a once-per-revolution reference

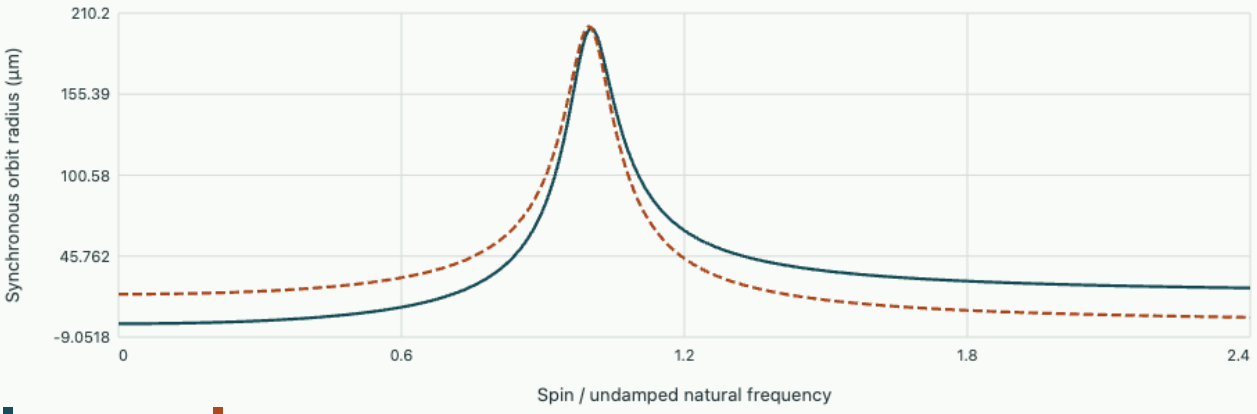
For $m = 2.0$ kg and $k = 80$ kN/m the undamped frequency is 200 rad/s or 1910 rpm. At $\zeta = 0.05, e = 20$ micrometres the shaft-centre orbit radius at $r = 1$ is 200 micrometres. A peak-to-peak probe reading across one axis would be 400 micrometres for a perfectly circular orbit, while an rms sinusoidal axis reading is 141 micrometres. State which amplitude is reported before comparing test and theory.

Open the live experiment

Try the calculation

Eccentricity (μm) **20**

Damping ratio **0.05**



At unit speed ratio the shaft-centre radius is 200.0 μm and its lag behind unbalance is 90°. The mass centre and shaft centre are different points.

The amplitude peak of a force proportional to speed squared occurs at $r = 1/\sqrt{1 - 2\zeta^2}$ when damping is sufficiently small, slightly above the undamped frequency. A constant-force resonance has its displacement peak slightly below. Near a fast run-up the steady expression need not apply because the rotor crosses the resonance faster than the transient can settle. Compare speed sweeps in both directions and hold the drive at representative speeds to distinguish sweep-rate memory from speed-dependent material stiffness.

Recover a complex influence coefficient from a trial mass

At a fixed operating speed write the measured synchronous probe response as a complex number relative to the once-per-turn mark. In linear operation a small calibrated trial unbalance adds a second complex response proportional to its mass-radius product. Subtract the baseline response as a vector; subtracting its magnitude would discard the phase information needed to identify the influence coefficient. The shaft centre model predicts a different phase as the test speed passes through resonance, so the coefficient belongs to the specific spin speed and instrument reference.

For an illustrative optical test the baseline response is 30 micrometres at zero reference phase. With the trial contribution present the response is 30 plus 20 imaginary micrometres in the same coordinates. Its displayed magnitude is 36.1 micrometres. The change of displayed magnitude is only 6.1 micrometres, while the actual trial response is 20 micrometres in quadrature. An estimate based on the two magnitudes would understate the response by more than a factor of three and lose its orientation. A convenient test of linearity is to repeat with the same trial pattern rotated through half a turn. To first order the complex change should reverse sign while the unchanged baseline remains.

When a probe sees a small geometric eccentricity of the actual shaft surface, its synchronous electrical reading includes this runout even at a very low speed where dynamic shaft bending is small. A slow-roll measurement recorded with the same phase mark provides an estimate of that vector. Subtract it in complex coordinates and test whether the resulting low-speed bending tends towards the square-of-speed prediction.

Tutorial · Infer damping and eccentricity separately

An isotropic test has known critical frequency. At exactly that speed the measured shaft radius is 150 micrometres; at high speed the radius tends to 15 micrometres. Estimate eccentricity and damping ratio. Why would one measurement at resonance be insufficient?

▼ **Use the two distinct asymptotes**

The high-speed magnitude tends to e, giving 15 micrometres if geometric surface runout has been removed. At $r = 1$, $Z/e = 1/(2\zeta)$; hence

$$\zeta = 15/(2 \cdot 150) = 0.050.$$

A resonance amplitude alone determines only the ratio e/ζ. High-speed data from a second bending mode or a noncircular shaft surface can violate the asymptote. An independently measured free-decay damping provides another way to separate the parameters without relying on a high-speed plateau.

Tutorial · Interpret orbit shape from two probes

Two calibrated orthogonal probes have synchronous amplitudes 100 and 60 micrometres with phase difference ninety degrees. What is the orbit and is that consistent with the simple isotropic response to one rotating unbalance?

▼ **Reconstruct the signed two-dimensional motion**

The trajectory is an ellipse with semiaxes 100 and 60 micrometres if those are peak amplitudes. Ideal isotropic stiffness and damping driven by a circular rotating imbalance give a circular synchronous orbit, so investigate support anisotropy, probe gains, cross-coupled bearings or another directional forcing. Swapping electronics and rotating probe orientation can separate an instrument gain difference from a physical ellipse. The sign of phase difference relative to the keyphasor determines whirl sense and cannot be recovered from the two amplitudes alone.

Tutorial · Locate a second peak without calling every peak critical

speed

A run-up shows one feature at a fixed shaft speed and a second one whenever a twice-per-turn component crosses the same structural frequency. How would an order–frequency map distinguish them and which terms are missing from the derived model?

▼ **Track spectral order and physical frequency**

The basic run-up forcing follows the one-per-turn line $f = \Omega/(2\pi)$; the second harmonic follows $2f$. If both excite the same fixed structural frequency at different speeds, their crossings align horizontally in physical-frequency coordinates but at separate shaft speeds. Noncircular runout, misalignment or other periodic effects can create higher shaft orders. The one-eccentric-disk equation contains only the synchronous term, so it cannot independently generate the second harmonic. Measure the forcing reference, examine phase and repeat after a calibrated trial mass. A changed one-per-turn response with a nearly unchanged twice-per-turn line helps distinguish the sources.

NPTEL Rotor Dynamics · Jeffcott rotor model and synchronous response

Vibration and identification in machinery · 2. Parametric resonance under periodic stiffness modulation

A carriage supported by a spring whose stiffness changes periodically can gain oscillation energy with no additive force at its natural frequency. Compare this mechanism with an ordinary forced resonance, estimate the first instability band and learn how to measure growth before nonlinear saturation invalidates the prediction.

A varying stiffness supplies work even at zero additive force

Let $k(t) = k_0[1 + h\cos(\Omega t)]$ with $|h| \ll 1$, and use a linear damper. The free coordinate obeys $\ddot{x} + 2\zeta\omega_0\dot{x} + \omega_0^2[1 + h\cos\Omega t]x = 0$, $\omega_0^2 = k_0/m$. Differentiating its instantaneous mechanical energy $m\dot{x}^2/2 + k(t)x^2/2$ and using the equation gives

$$\dot{E} = -c\dot{x}^2 + \dot{k}x^2/2.$$

Modulating stiffness can therefore do positive or negative work depending on the oscillation phase. For an exactly zero initial displacement and velocity the ideal equation stays at zero; real noise or a deliberate small initial motion supplies a seed.

The squared displacement has a component at twice the oscillation frequency. A stiffness variation near twice ω_0 can preferentially raise stiffness at large displacement and lower it near a zero crossing, supplying net energy over a cycle. This identifies the principal parametric band but does not determine its threshold. A nonzero damping loss must first be overcome.

Background · At what point in a cycle should the stiffness increase to feed motion?

Imagine a spring whose stiffness can switch between two nearby values fast enough that displacement and velocity stay continuous during a switch. The instantaneous change in mechanical energy equals one half of the stiffness change times displacement squared. Increase stiffness at the two displacement extrema and decrease it at the two zero crossings. For an almost sinusoidal motion of amplitude A, the leading injected energy per cycle is

$$\Delta E_{\mathrm{pump}} \simeq 2\left(\tfrac{1}{2}\Delta k\, A^2\right) - 2\left(\tfrac{1}{2}\Delta k\, 0^2\right) = \Delta k\, A^2.$$

Reverse the switching phase and the net transfer changes sign. This accounting explains both the importance of phase and the two stiffness changes of the same sign per mechanical period. The ideal switches describe a square modulation; the numerical threshold for sinusoidal modulation must instead use its own Fourier amplitude and the damping loss.

Find slow growth and detuning near twice resonance

Write the motion as slowly varying cosine and sine quadratures at half the modulation frequency. Substitute into the oscillator equation, drop second derivatives of slow amplitudes and average out

terms oscillating near three times the basic frequency. With detuning $\Delta = \Omega/2 - \omega_0$ and retaining first orders in small damping, modulation and detuning, the two coupled quadratures have exponents

$$s_{\pm} = -\zeta\omega_0 \pm \sqrt{(h\omega_0/4)^2 - \Delta^2}.$$

The square-root term is real sufficiently near the principal band. Net growth requires its positive branch to exceed damping. At exact tuning the approximate threshold is $h > 4\zeta$; away from tuning the unstable interval narrows according to $\Delta^2 < (h\omega_0/4)^2 - (\zeta\omega_0)^2$. The calculation assumes the damping, relative modulation and fractional detuning are small. For large modulation use the periodic linear system and compute its Floquet multipliers directly.

For $\omega_0 = 100$ rad/s, $\zeta = 0.015$ and $h = 0.10$, exact-tuning growth is $2.5 - 1.5 = 1.0$ per second. The unstable half-width in $\Omega/2$ is 2.0 rad/s, hence in the actual drive Ω it is 4.0 rad/s. A one-percent change of pump frequency near 200 rad/s already changes half-frequency detuning by 1 rad/s. Distinguish the drive-frequency width from oscillation-frequency width when comparing an experimental scan.

Compare this finite instability band with the bounded steady response of a damped rotating force

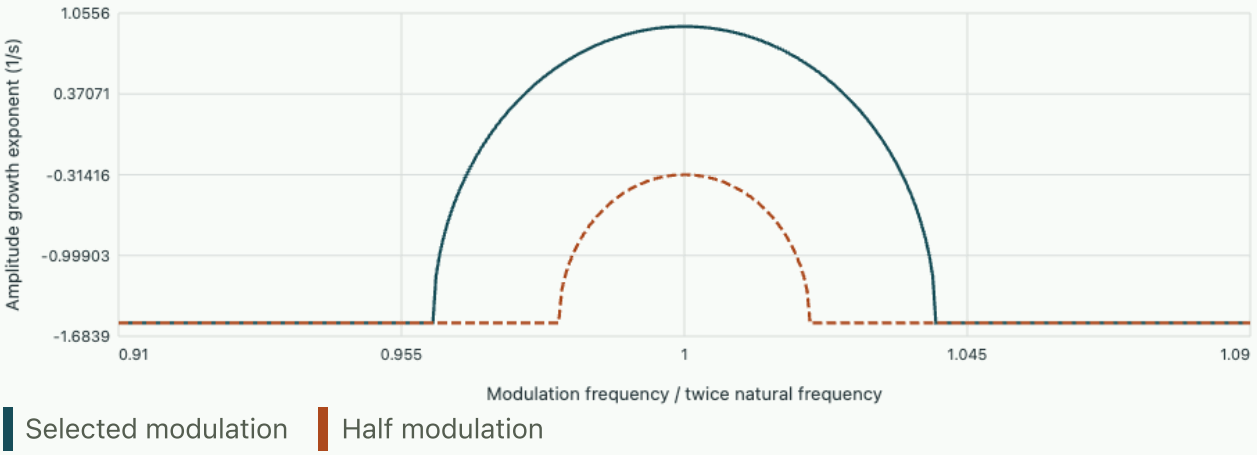
Distinguish the real part of a growth exponent from detuning phase

Outside the principal resonance tongue the square root in the approximate exponent becomes imaginary. The oscillation phase then drifts relative to half the modulation frequency, while the real part remains the viscous decay rate at this order. A graph showing only the real part therefore develops a flat negative floor. This floor means exponential decay of the envelope, not that the instantaneous displacement has stopped oscillating. At the tongue boundaries the larger real exponent crosses zero; inside the band the early envelope increases if the initial condition contains the growing quadrature.

Open the live experiment

Try the calculation

Natural frequency (Hz)	**10**
Stiffness modulation fraction	**0.16**
Damping ratio	**0.025**



The leading-order exact-tuning exponent is 0.942 per second. The predicted unstable band is 19.38 to 20.62 Hz.

At a natural frequency of ten hertz, modulation fraction 0.16 and damping ratio 0.025, the exact-tuning early exponent is about 0.942 per second. An amplitude doubles in approximately 0.736 seconds. Halving the modulation gives a negative exponent at the centre, even though the stiffness is still being varied periodically. Compare measured amplitude over many oscillation periods but stay within the early linear regime. Starting the drive with a phase almost entirely in the initially decaying quadrature can delay the emergence of the growing component and make a short time record look stable.

A direct numerical check uses a pair of independent initial states. Integrate the original periodic oscillator for one modulation period from each initial state, then place their final states in the columns of a two-by-two transition matrix. Repeating the evolution multiplies by this same matrix. A multiplier of magnitude greater than one means that a suitable initial state grows from period to period. The logarithm of its magnitude divided by the period estimates a growth rate without assuming weak detuning. Agreement with the simpler displayed equation should improve near the primary band when damping and modulation both decrease.

Tutorial · Determine the threshold from a measured decay

With modulation off, a freely released oscillator decays in amplitude as $e^{-0.80t}$ and has natural frequency 80 rad/s. Estimate the damping ratio and minimum small periodic stiffness variation required near exact principal tuning.

▼ **Translate an amplitude decay into a pump threshold**

The measured decay rate is $\zeta\omega_0$, so $\zeta = 0.010$ and the threshold is about $h = 0.040$, a four-percent relative modulation amplitude. If the measured quantity were energy rather than amplitude, its exponential rate would be twice as large for the same damping. A stiffness swing quoted peak-to-peak is also twice the h amplitude in the stated equation. Both conventions must be resolved before claiming agreement to a factor of two.

Tutorial · Predict whether a detuned pump grows

For the numerical 100 rad/s oscillator driven with $h = 0.10$, compare modulation frequencies 202 and 205 rad/s. Estimate growth or decay of the dominant quadrature.

▼ **Use half the drive detuning**

At 202, $\Delta = 1$, giving $s_+ = -1.5 + \sqrt{6.25 - 1} = 0.791$ per second, hence growth. At 205, $\Delta = 2.5$ and the leading square root vanishes, giving decay at 1.5 per second at this approximation level. Starting in a predominantly decaying quadrature can delay visibly growing motion even inside the band. Record a long enough early linear window and vary initial phase rather than inferring stability from one quiet short trace.

Tutorial · Use a one-period map beyond weak modulation

Describe a numerical procedure that determines linear stability without fitting an arbitrary initial envelope. How can it be checked in the limit of zero periodic modulation?

▼ **Propagate independent displacement and velocity states**

Integrate two independent initial states, for example unit displacement with zero velocity and vice versa, through one pump period. Place their final states in the columns of a monodromy matrix. An eigenvalue of magnitude greater than one implies linear growth with repeated periods. Its magnitude sets the growth rate through its logarithm divided by the period. Refine the integration time step until both multipliers converge. The determinant should be $e^{-2\zeta\omega_0 T}$ from the trace of the state matrix. With $h = 0$ and underdamping both multipliers have magnitude $e^{-\zeta\omega_0 T}$. A nonlinear spring or motion stops ultimately limit real growing amplitudes; linear multipliers predict onset and early growth, not the final oscillation level.

NPTEL Mechanical Vibrations · parametrically excited systems and stability

Vibration and identification in machinery · 3. Random base motion, spectral density and rms response

A support vibration is described by a power spectral density rather than a single sine-wave amplitude. Derive the relative-motion filter for a mounted instrument, calculate mean-square motion from a measured spectrum and distinguish a stationary rms prediction from a peak displacement guarantee.

Start from moving-base force balance

Let $y(t)$ be base displacement, $x(t)$ the instrument's absolute displacement and $z = x - y$ relative travel inside its mount. Newton's law for the mass is $m\ddot{x} = -c(\dot{x} - \dot{y}) - k(x - y)$. Therefore

$$\ddot{z} + 2\zeta\omega_n\dot{z} + \omega_n^2 z = -\ddot{y}.$$

In the frequency domain the transfer magnitude squared from measured base acceleration to relative travel is

$$|H_{za}(\omega)|^2 = \frac{1}{(\omega_n^2 - \omega^2)^2 + (2\zeta\omega_n\omega)^2}.$$

A base accelerometer's physical units are metres per second squared. If it reports in multiples of g, convert amplitude by 9.81 and spectral density by 9.81^2. A displacement spectrum is not obtained simply by treating base acceleration as an applied displacement; the numerator and units differ. At sufficiently low frequency a constant base acceleration produces a quasi-static relative deflection $-a/\omega_n^2$, which checks the zero-frequency transfer.

A variance is an area over a stated frequency convention

For a stationary linear response and one-sided base acceleration spectrum $S_a(f)$ per hertz, the relative mean-square motion is $\int_0^\infty |H_{za}(2\pi f)|^2 S_a(f)\,df$. The integral treats distinct frequency components as uncorrelated under stationarity. If a two-sided density per radian per second is used instead, both frequency integration and normalisation must change consistently. The integral of an input spectrum over its own stated axis must reproduce the independently measured acceleration variance.

A convenient benchmark assumes a constant one-sided spectrum S_0 across all frequencies relevant to the mode. Integrating the squared response gives

$$\langle z^2 \rangle = \frac{S_0}{8\zeta\omega_n^3}.$$

With $\omega_n = 100$ rad/s, $\zeta = 0.05$ and $S_0 = 0.040\ (\mathrm{m/s^2})^2/\mathrm{Hz}$, rms relative motion is 0.316 mm. A real finite passband and nonwhite input should be integrated from its measured spectrum; extending a nominal constant outside its measured range is defensible only if the omitted contribution is small. An ideal white acceleration input over all frequencies has infinite acceleration variance, yet this relative displacement filter yields a finite response variance. It serves as a response benchmark, not a physically realisable unlimited-band source.

Compare relative mount travel with the harmonic isolation and transmissibility chapter

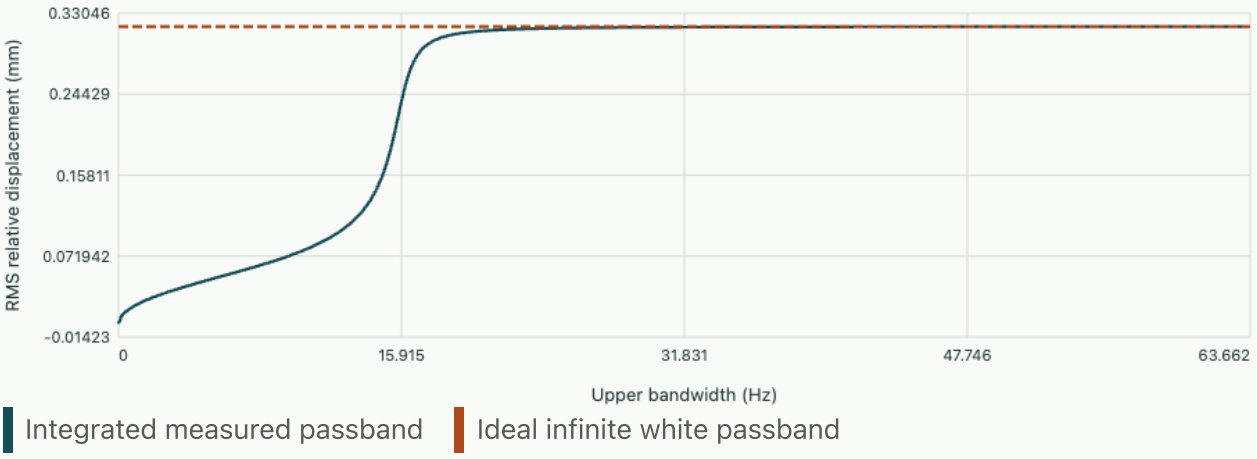
Integrate the spectrum before reporting a single motion amplitude

The experiment accumulates output variance from zero frequency up to a movable upper bandwidth. Most of the increase occurs close to the mode when damping is small. Below that region the support acceleration can contain measurable energy but the relative mount travel stays small; far above resonance the displacement response to acceleration drops quickly. Read the asymptote as a numerical integration check. A response moving toward the predicted limit does not imply that a real excitation spectrum remains flat beyond the accelerometer bandwidth.

Open the live experiment

Try the calculation

Natural angular frequency (rad/s)	100
Damping ratio	0.05
One-sided acceleration PSD ((m/s²)²/Hz)	0.04



The ideal white-input limit is 0.316 mm rms; stopping the integration at 63.7 Hz gives 0.316 mm.

With a natural frequency near sixteen hertz and damping ratio 0.05, the modal half-power width is roughly 1.6 hertz. A one-second record gives frequency spacing around one hertz, leaving very few independent locations across the resonance. A longer record improves frequency resolution. Dividing the record into many shorter segments and averaging their estimates reduces scatter while broadening the resolvable features. State both record duration and segment length and inspect how changing them affects the estimated cumulative variance.

A narrow deterministic tone superposed on random acceleration requires care. Its line energy is distributed across nearby bins according to the window; a density value at the tallest bin changes with resolution even when the physical tone amplitude does not. Integrate the line over its full width or fit a

sinusoid in the time domain. Then combine its mean-square output with the background variance if the two contributions have zero cross correlation. For example, a statistically independent 0.20 mm rms line and 0.30 mm rms broadband contribution produce about 0.361 mm rms total. Adding the two rms amplitudes would yield 0.50 mm, corresponding to a very different assumption about phase and dependence.

A final consistency check uses a short time-domain record with an independently measured travel sensor. Remove a known static offset, compute its root mean square and compare with the integral over the sensor's actual bandwidth. Compare the same frequency interval in both records before attributing a discrepancy to modal damping.

Tutorial · Verify a finite-band low-frequency calculation

A base spectrum is flat at 0.020 acceleration-squared per hertz between 0 and 2 Hz, negligible elsewhere. For a mount with $\omega_n = 200$ rad/s, estimate base rms acceleration and relative displacement rms. Why does the white-through-resonance formula fail here?

▼ Apply the low-frequency transfer over the measured band

Base variance is $0.020(2) = 0.040$, giving rms 0.20 m/s². All excited frequencies are well below the mount resonance near 31.8 Hz, so rms relative travel is approximately $0.20/200^2 = 5.0$ micrometres. The broadband result would assume input energy near the resonant amplification peak that this spectrum explicitly lacks. Direct numerical quadrature with the exact transfer checks the small low-frequency approximation.

Tutorial · Distinguish rms from a time-dependent peak

A zero-mean Gaussian displacement record has the 0.316 mm rms value of the broadband example. What does a three-rms clearance mean at a single time, and why does it not provide that same probability of never touching the stop over a long run?

▼ Include time correlation and observation duration

At a random stationary time a Gaussian value lies outside plus or minus three standard deviations with probability about 0.0027. Over many samples or many correlation times there are repeated opportunities for an excursion. Independent observations would give probability $1 - (1 - 0.0027)^N$ of at least one exceedance, whereas actual filtered vibration has correlated values. Peak predictions need the observation duration and a model of frequency content or crossings. Mechanical stops introduce nonlinearity after contact and invalidate the original Gaussian linear response thereafter.

Tutorial · Identify insufficient frequency resolution

A lightly damped resonance at 20 Hz has damping ratio 0.01. A measured transfer and input spectrum are tabulated every 1 Hz. Estimate its half-power bandwidth and design a quadrature and measurement check.

▼ **Resolve the resonance before integrating its area**

For a lightly damped isolated mode the full bandwidth is approximately $2\zeta f_n = 0.40$ Hz. A one-hertz grid can miss most of the peak or place a single anomalously high bin at its centre. Use a much finer frequency step or fit a justified complex modal response, then integrate the product with the measured input spectrum. Increase acquisition duration to obtain true frequency resolution; interpolating sparse measured bins alone cannot reveal an unmeasured narrow input line. Compare predicted rms with a time-domain record and account for window leakage and accelerometer noise.

NPTEL Elements of Mechanical Vibration · random vibration, spectral density and system response

Vibration and identification in machinery · 4. Why a local damper couples otherwise separate modes

Undamped mode shapes diagonalise a symmetric mass and stiffness model. Adding a damper at one attachment does not generally leave modal equations separate. Derive the resulting coupling, test its magnitude and learn when a diagonal modal-damping approximation can hide an observable phase error.

Transform the entire equation, including the damper

For two degrees of freedom write $M\ddot{x} + C\dot{x} + Kx = f$, with symmetric positive definite M, K and passive nonnegative C. Choose real undamped eigenvectors as columns of Φ, mass normalised so $\Phi^{\mathsf{T}} M \Phi = I$ and $\Phi^{\mathsf{T}} K \Phi = \operatorname{diag}(\omega_1^2, \omega_2^2)$. Setting $x = \Phi q$ gives transformed damping $C_m = \Phi^{\mathsf{T}} C \Phi$. This matrix is generally full. Orthogonality established for mass and stiffness is not automatically orthogonality for an independently located dissipative element.

A Rayleigh model $C = aM + bK$ is a sufficient condition for diagonal modal damping, with coefficients $a + b\omega_j^2$. It is not necessary in every multimode setting, nor does fitting two measured damping ratios prove the actual physical damping follows this law across other modes. An isolated dashpot to ground at coordinate one produces $C = c e_1 e_1^{\mathsf{T}}$, hence $(C_m)_{jk} = c\phi_{1j}\phi_{1k}$. If both modes move appreciably at that attachment, off-diagonal damping is unavoidable.

Compute a forced response in both modal coordinates

Let $M = I$ kg and let symmetric and antisymmetric mass-normalised vectors $(1,1)/\sqrt{2}$, $(1,-1)/\sqrt{2}$ have natural frequencies 10 and 14 rad/s. A dashpot of 4 N s/m on coordinate one gives modal damping matrix with every entry 2 per second. At forcing frequency ω, the complex dynamic stiffness has diagonal entries $\omega_j^2 - \omega^2 + 2i\omega$ and off-diagonal entries $2i\omega$. Apply a modal force only in the first coordinate. The second equation now requires

$$\frac{Q_2}{Q_1} = -\frac{2i\omega}{\omega_2^2 - \omega^2 + 2i\omega}.$$

At $\omega = 10$ rad/s its magnitude is

$$20/\sqrt{96^2 + 20^2} = 0.204.$$

A diagonal modal-damping approximation would predict exactly zero second modal response. Both models dissipate energy and can give similar-width first-mode peaks, yet their reconstructed physical coordinates differ in phase and amplitude. At larger mode spacing or lower damping the off-diagonal transfer becomes smaller, supplying a quantitative rather than categorical justification for ignoring it.

Compare a single circular harmonic response with an instrument observing two coupled modal coordinates

Check passivity independently of fitted modal damping ratios

With external force zero the mechanical energy derivative is $-\dot{x}^{\mathsf{T}} C \dot{x}$. For the single grounded dashpot this is $-4\dot{x}_1^2 \leq 0$. In modal coordinates it is $-2(\dot{q}_1 + \dot{q}_2)^2$, including a cross term. Dropping that term assigns different instantaneous dissipation to particular velocity combinations. When $\dot{q}_1 = -\dot{q}_2$, the physical damped coordinate is momentarily stationary and true instantaneous loss is zero; a diagonal approximation incorrectly predicts positive loss. Subsequent forces can move that coordinate, so a lossless instant does not imply an undamped eigenmode.

Interpret the second modal coordinate in an experiment

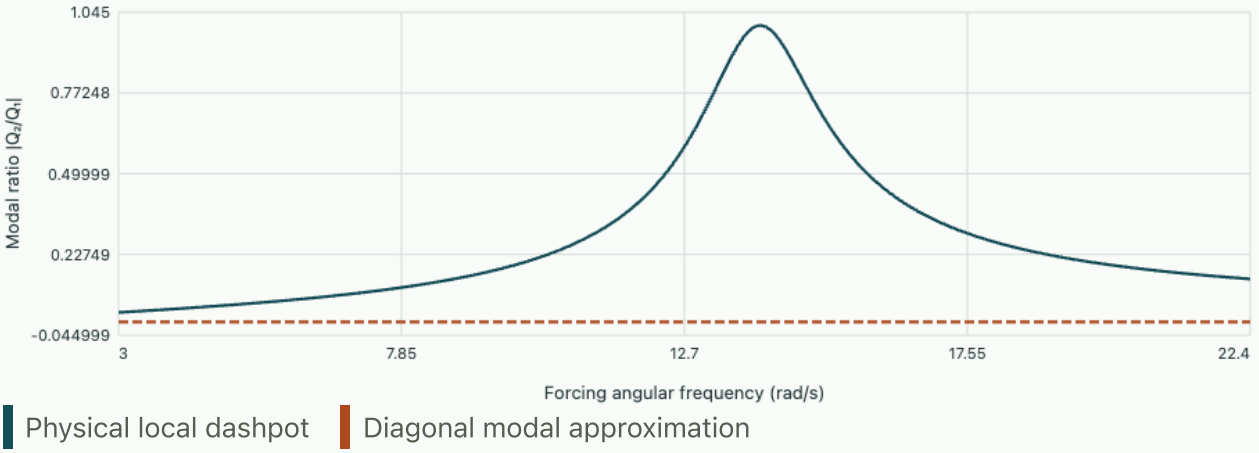
A laboratory cannot normally measure the modal coordinates directly. For the symmetric two-mass example, mount matched accelerometers on the two masses and compute sum and difference using a common clock. For steady harmonic response both measured accelerations share the factor of minus frequency squared relative to their displacements. The ratio of the difference channel to the sum channel therefore recovers the same modal displacement ratio, provided the sensors use the same sign convention and gain. A unit gain error or a cable polarity reversal changes the inferred cross-response even if the physical system is stationary.

Open the live experiment

Try the calculation

Local dashpot (N s/m) **4**

Second modal frequency (rad/s) **14**



At the first natural frequency of 10 rad/s, the second-to-first modal amplitude ratio is 0.204. At the undamped second frequency the ratio is 1.00.

Apply equal-phase forces to the two masses using a known load-sharing beam. That forcing is purely symmetric in the ideal model. An independently measured small antisymmetric force contaminates the nominally unforced modal coordinate. Shift the local damper to the opposite mass and repeat at the

same frequency. The off-diagonal terms of the modal damping matrix reverse sign while the diagonal terms do not. The induced antisymmetric response should reverse phase relative to the symmetric one; an unchanged spurious force is less likely to reproduce that reversal. The sign information is absent from the displayed ratio magnitude and must be preserved in the experimental phase records.

A useful limiting case sets the forcing frequency equal to the undamped second natural frequency. Its real diagonal stiffness then vanishes and the local-damping term directly balances the coupling from the first modal coordinate. In this example their ratio has unit magnitude. A large ratio can occur even when the applied modal force has no second component. This observation follows from the coupled equations at that frequency and does not imply a universal mode shape or a one-to-one transfer of energy between modes.

Tutorial · Find a damping placement that restores modal separation

Keep the same grounded dashpot on coordinate one and add another identical dashpot on coordinate two. Compute the new modal matrix and dissipation. What happens if the added dashpot is only half as large?

▼ Add physical contributions before transforming

Equal grounded dampers give $C = 4I$, so $C_m = 4I$. Each mode then has amplitude damping rate two per second and no damping coupling. If the second is only 2 N s/m, C_m has diagonal entries three and off-diagonal entries one. Common and differential combinations have different physical loss. Symmetric placement removes the off-diagonal term because the two undamped shapes have opposite relative signs on the second coordinate.

Tutorial · Reconstruct both modes from two physical probes

At one frequency the calibrated complex sensor readings are $X_1 = 1.0$ mm and $X_2 = 0.8 + 0.2i$ mm. Determine Q_1, Q_2 and assess whether a second mode is detectable even though no second peak has been observed.

▼ Transform complex amplitudes without discarding phase

In this normalisation $Q_1 = (X_1 + X_2)/\sqrt{2} = (1.8 + 0.2i)/\sqrt{2}$ and $Q_2 = (X_1 - X_2)/\sqrt{2} = (0.2 - 0.2i)/\sqrt{2}$ mm in the corresponding modal coordinate units. The magnitude ratio is approximately 0.156. A single-magnitude sensor cannot perform this separation, and two magnitude-only readings lose the relative phase. Compare this ratio with a noise and gain-mismatch estimate before declaring it physical. A weakly present mode need not appear as a visually isolated peak at the measured driving frequency.

Tutorial · Determine the true damped poles

How should the complete free-decay problem be solved when modal damping is non-diagonal? Give a diagnostic distinguishing a fitted damping error from an incorrectly specified mass or stiffness.

▼ Use a state-space or quadratic eigenproblem

Seek $x = ve^{st}$, obtaining $(s^2M + sC + K)v = 0$. Solve the quadratic eigenproblem or convert it to first-order state space with displacement and velocity. Damped eigenvectors can be complex even though all physical matrices are real, and complex poles occur in conjugate pairs. Compare both measured phases and decay rates at multiple sensors, then repeat after altering the dashpot while leaving masses and elastic supports unchanged. Frequency shifts or mode-shape changes present already with the damper removed suggest an underlying mass or stiffness error. A fit to one resonance width alone cannot identify the full matrix.

Adhikari · Optimal complex modes and an index of damping non-proportionality, Cambridge research paper

Vibration and identification in machinery · 5. Which damping law produced the free-decay trace?

A carriage slides on a guide and a spring pulls it back to its resting position. Two dampers can be adjusted to give the same first peak reduction. Predict the following turning points and decide which measurement separates viscous loss from dry sliding friction.

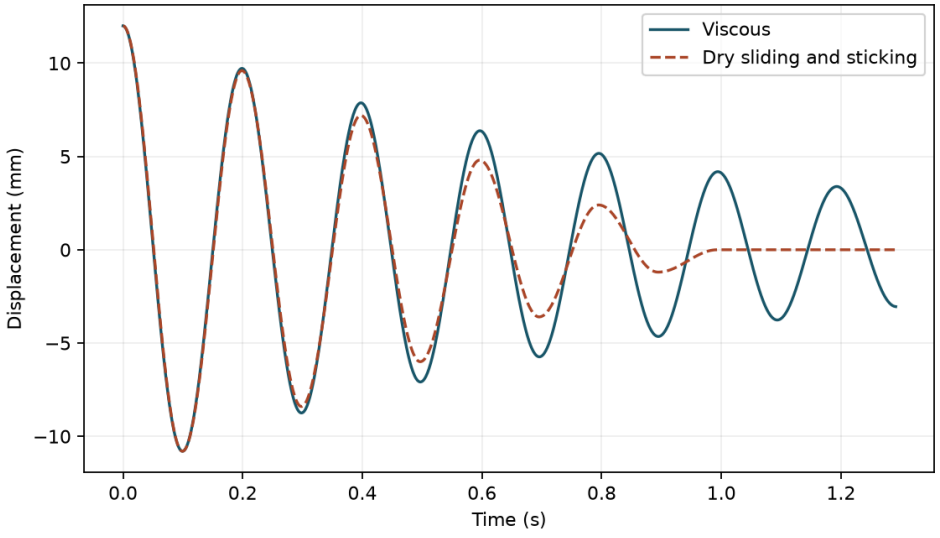
Use the turning points to identify the energy loss mechanism

For a viscous damper the free equation is $m\ddot{x} + c\dot{x} + kx = 0$. In the underdamped regime let $\omega_n = \sqrt{k/m}$, $\zeta = c/(2\sqrt{km})$ and $\omega_d = \omega_n\sqrt{1-\zeta^2}$. Released from $x(0) = A_0$ with zero velocity, the displacement is $A_0 e^{-\zeta\omega_n t}[\cos\omega_d t + (\zeta\omega_n/\omega_d)\sin\omega_d t]$. Turning points occur every π/ω_d. Their absolute sizes have a constant ratio $e^{-\pi\zeta/\sqrt{1-\zeta^2}}$.

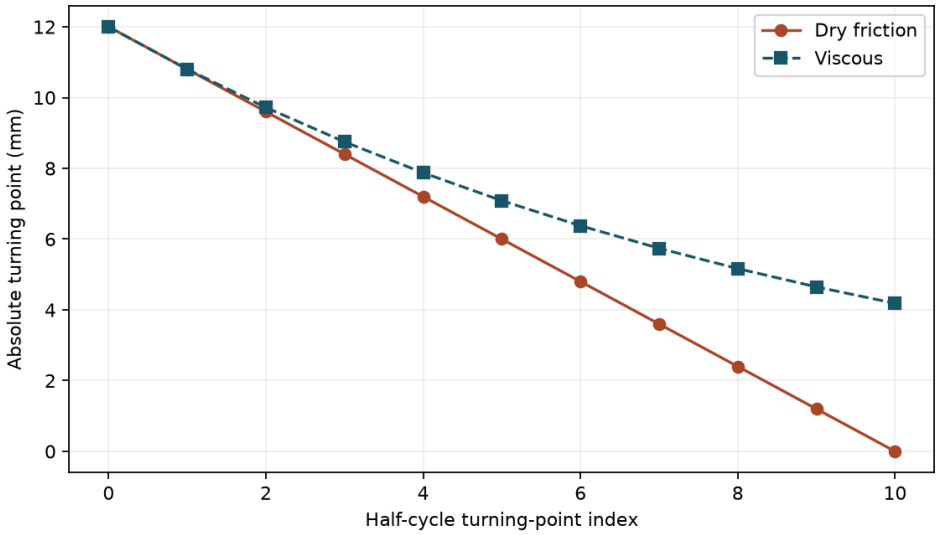
Now let a dry guide exert a constant opposing force of magnitude F. During a sliding half-cycle its sign remains fixed. Starting from a positive peak with $kx > F$, the spring accelerates the mass to the left while friction acts to the right. Subtract the shifted equilibrium F/k from the coordinate. The remainder obeys an undamped oscillator equation, so the next turning point after π/ω_n is $2F/k - x$. For a release from the other side, replace F/k by its signed counterpart. The magnitudes of successive peaks fall by $2F/k$ while the mass still crosses the origin. At a turning point where $|kx| \le F$, static friction can support the spring and the body sticks there. Here the static threshold is idealised as equal to the sliding force.

Match the first peak and predict the next three

Let $m = 1$ kg, $k = 1000$ N/m, $F = 0.6$ N and $A_0 = 12$ mm. Dry friction subtracts 1.2 mm of magnitude per sliding half-cycle. The first return is at –10.8 mm. A viscous model can reproduce its magnitude if its half-period peak ratio is 0.9. Solving $\delta = \pi\zeta/\sqrt{1-\zeta^2}$ for $\delta = -\ln(0.9)$ gives $\zeta = 0.0335$. Its fourth half-cycle peak is $12(0.9)^4 = 7.87$ mm, whereas the dry-friction prediction is 7.20 mm. A single peak pair cannot distinguish the models. Several peaks well above the instrument resolution can.



The sliding model is evaluated on exact half-cycle segments. A plateau is physical when the static force threshold can hold the carriage.



Plotting amplitude against peak count makes an almost missed difference in the first displacement trace easier to measure.

Exercise · estimate sliding friction

A 1000 N/m guide returns through peak magnitudes 7.2, 6.0 and 4.8 mm. Estimate its friction and predict the following magnitude.

▼ **Work through the solution**

The decrement per half-cycle is 1.2 mm. The corresponding force is $k(0.0012)/2 = 0.60$ N. The next alternating peak magnitude is 3.6 mm, still well above the sticking threshold of 0.6 mm.

Exercise · check the last interval

A positive turning point is 0.9 mm and $F/k = 0.6$ mm. Will the mass reach a negative turning point?

▼ **Work through the solution**

Sliding restarts because $0.9 > 0.6$. The next position is $2(0.6) - 0.9 = 0.3$ mm, still positive. Since $0.3 < 0.6$, static friction can hold it there. Extrapolating a positive linear amplitude envelope through zero would lose the actual final position.

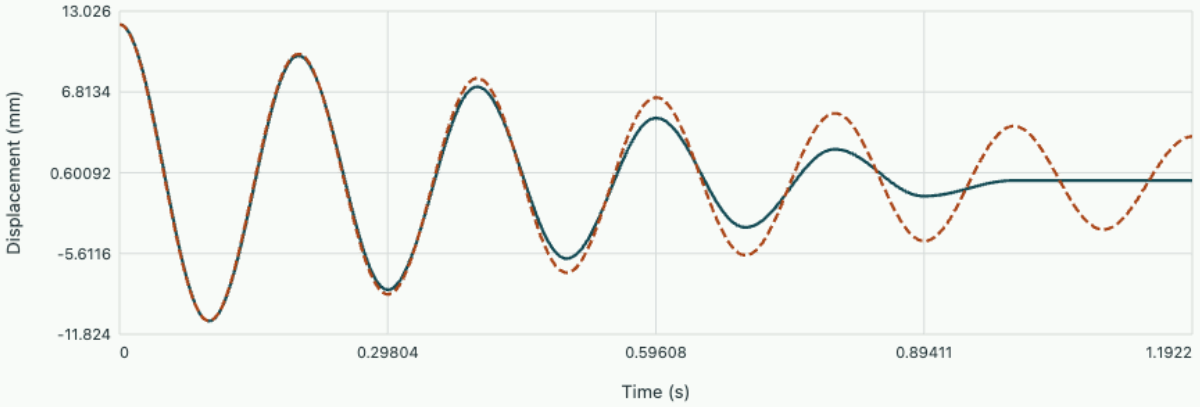
Interactive calculation · Dry vs viscous

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

Release amplitude (mm)	**12**
Sliding friction (N)	**0.6**
Stiffness (N/m)	**1000**



Coulomb sliding and final stick Viscous, matched first turning point

Matched viscous damping ratio 0.0335. Dry friction stops after 10 half cycles at 0.000 mm; the viscous model approaches zero asymptotically. Mass1 kg.

Download this example's editable notebook

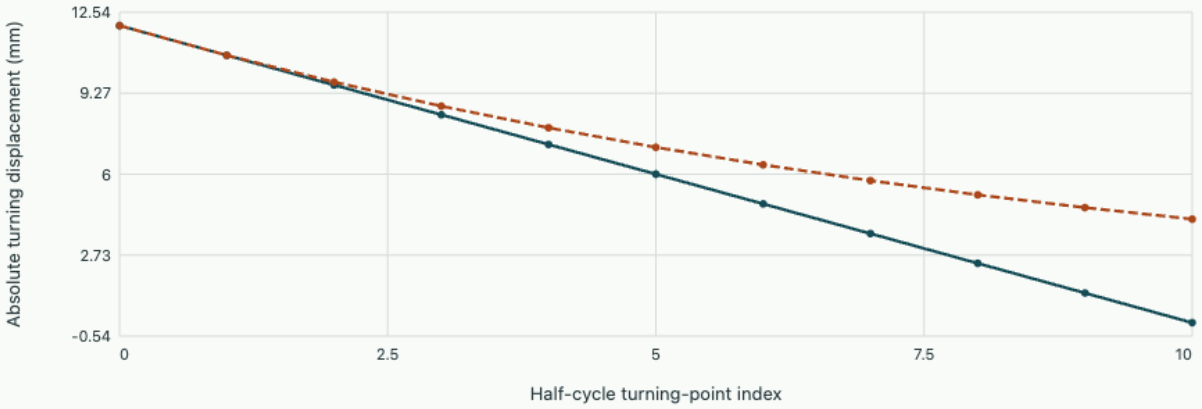
Interactive calculation · constant loss versus fractional loss

Compare turning points from the same piecewise sliding solution and matched viscous oscillator. Continue until dry friction reaches its sticking condition.

Open this live calculation

Try the calculation

Release amplitude (mm)	12
Sliding friction (N)	0.6
Stiffness (N/m)	1000



Coulomb, terminating in stick Matched viscous decay

Dry sliding loses 1.200 mm in turning amplitude per complete half cycle, until sticking at 0.000 mm. The viscous sequence loses a fixed fraction 0.100 instead. Its index counts its own slightly longer damped half cycles.

Download the editable notebook and compare both observables

Open editable notebook Download .ipynb

Vibration and identification in machinery · 6. From beam curvature to a mass matrix and a frequency measurement

A coarse cantilever model often predicts its first resonance accurately while putting the third far too high. Derive both matrices from the same displacement field and inspect how each mode converges.

Turn distributed kinetic and strain energy into coordinates

Take a uniform slender beam at rest, with transverse displacement $w(x,t)$ and small rotation. Neglect cross-section rotary inertia, shear compliance and prestress. An element of length h stores strain energy $\int EI(w'')^2 dx/2$ and kinetic energy $\int \rho A \dot{w}^2 dx/2$. Continuity of both displacement and slope at the interfaces prevents a spurious concentrated curvature. A cubic polynomial has four coefficients, exactly enough to interpolate the displacement and slope at each end.

Put $s = x/h$ and $q_e = (w_1, \theta_1, w_2, \theta_2)^T$. Apply the four endpoint conditions to a cubic, or verify them directly, to obtain $w = N q_e$ with

$$N = (1 - 3s^2 + 2s^3,\ h(s - 2s^2 + s^3),\ 3s^2 - 2s^3,\ h(-s^2 + s^3)).$$

Substitution into the energy integrals gives the element stiffness and consistent mass. The word consistent records that the same shape functions transport mass and displacement.

$$K_e = \int_0^h EI (N'')^T N'' \, dx,$$

$$M_e = \int_0^h \rho A N^T N \, dx.$$

For example, $N_1'' = (-6 + 12s)/h^2$, so $K_{11} = 12EI/h^3$. Integrating N_1^2 gives $M_{11} = 156\rho A h/420$. Mass couplings to rotations carry one additional length factor, and rotation–rotation terms carry two. These dimensions catch sign and assembly errors which a plausible first eigenvalue can hide.

Recover the vibration equation and its boundary conditions

Scatter-add each element into both global matrices at its shared nodal indices. A clamped root imposes $w_0 = \theta_0 = 0$. Delete those rows and columns from both matrices. The free end has no applied transverse force or moment, which the weak form already enforces as natural conditions. Lagrange's equations give $M\ddot{q} + Kq = 0$. A harmonic ansatz gives $K\phi = \omega^2 M \phi$. Eigenvectors from different frequencies are orthogonal in the mass inner product. Their overall signs remain arbitrary.

Take a one-metre steel cantilever, 30 mm wide and 10 mm deep, $E = 210$ GPa and $\rho = 7800$ kg/m³. The analytical clamped-free roots solve $\cosh\beta \cos\beta = -1$. Their frequencies scale as $\beta^2 \sqrt{EI/(\rho A)}/(2\pi L^2)$. The test provides three independent references and a direct check on units and boundary conditions.

Interactive experiment · Rotate a beam mode and vary length, thickness and phase

The live plot evaluates the analytical clamped-free mode. The spatial view gives the beam its rectangular cross-section so that the modal centreline can be inspected from different directions. Move the phase slider through one cycle; the displayed frequency gives the physical time corresponding to that cycle.

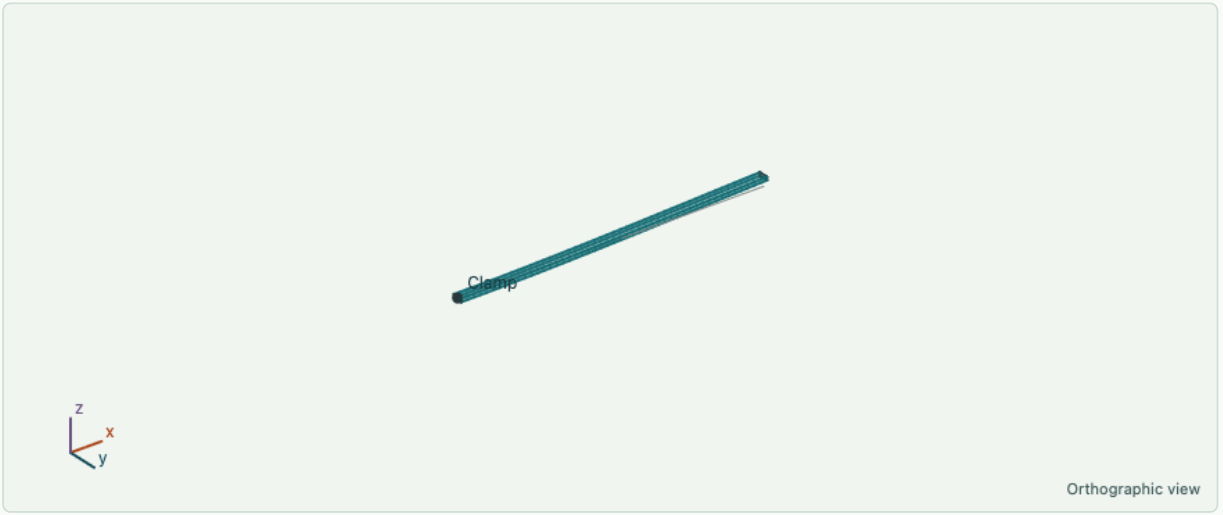
$$\phi(x) = \cosh(\beta x/L) - \cos(\beta x/L) - \frac{\cosh\beta + \cos\beta}{\sinh\beta + \sin\beta}\left[\sinh(\beta x/L) - \sin(\beta x/L)\right].$$

Scale this shape to the chosen tip amplitude and multiply it by the cosine of the phase. Double the length and check the factor of four in period; double the depth and check the factor of two in frequency. Modes two and three introduce internal nodes which a coarse mesh may miss. The interactive view uses the analytical slender-beam solution; the OneSimulation notebook below independently assembles and solves the discrete mass–stiffness eigenproblem. The selected amplitude illustrates a linear mode and is not a load-response prediction.

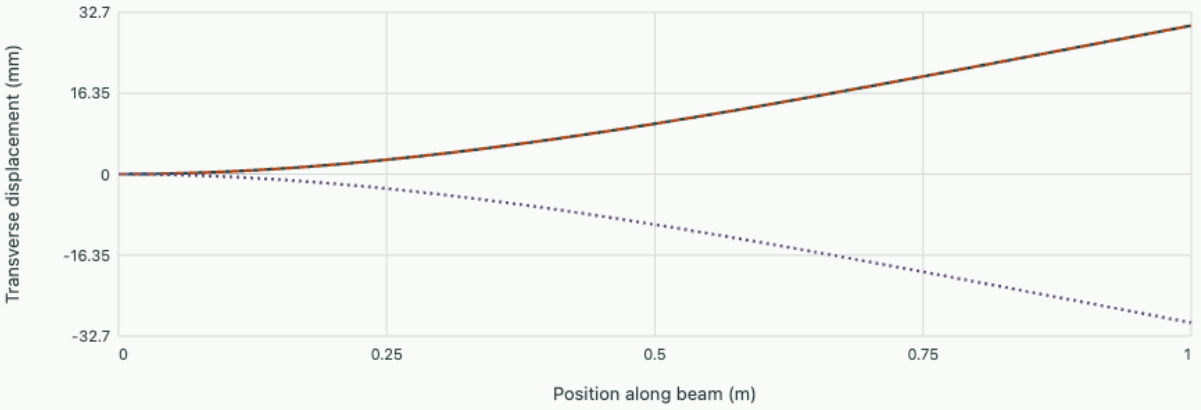
Open this experiment online

Try the calculation

Mode number	1
Beam length (m)	1
Beam depth (mm)	10
Tip amplitude (mm)	30
Oscillation phase (degrees)	0



A three-dimensional extrusion of the analytical beam mode. Width is 30 mm; depth and modal phase are adjustable. The section is drawn without the small rotation correction. Frequency and centreline follow the slender-beam model; the companion notebook checks the numerical eigenproblem. Axes: x = x (m); y = y (m); z = z (m).



Selected instant | Positive modal envelope | Negative modal envelope

Mode 1: 8.3819 Hz; period 119.305 ms. The analytical Euler–Bernoulli mode has a clamped root and a free tip. Shape amplitude is an illustration of the linear mode, not a predicted forced response.

Continue in the OneSimulation notebook · Download its Python model

OneSimulation experiment · Convergence of the first three cantilever frequencies under mesh refinement

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

```
"""Assemble a cantilever from Hermite elements; validate its first three modes."""
import numpy as np
import matplotlib.pyplot as plt
from onesim.numerics.linear import SparseSymmetricEigenSolver

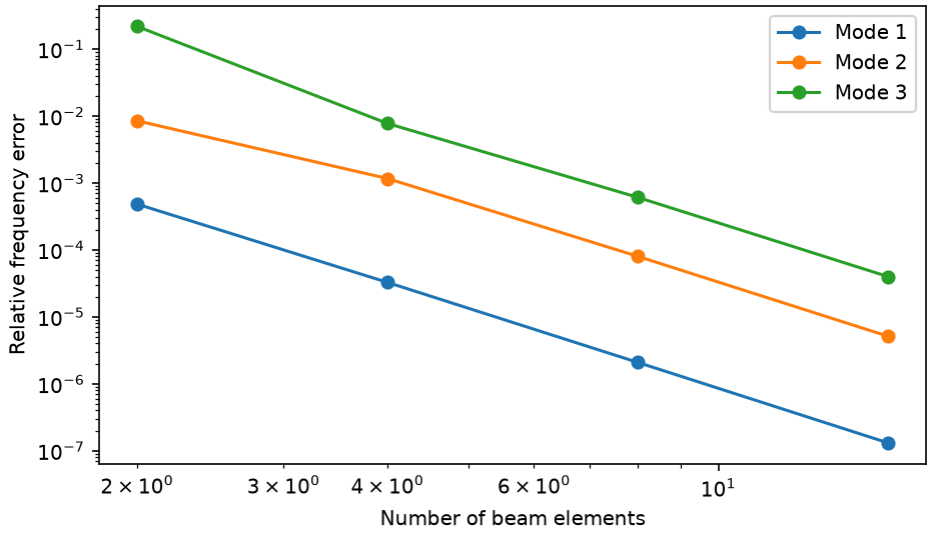
length, width, depth = 1.0, 0.03, 0.01  # m
E, density = 210e9, 7800.0
area, inertia = width*depth, width*depth**3/12

def beam_frequencies(elements):
    h = length/elements
    k = E*inertia/h**3*np.array([[12,6*h,-12,6*h],[6*h,4*h*h,-6*h,2*h*h],
                                 [-12,-6*h,12,-6*h],[6*h,2*h*h,-6*h,4*h*h]])
    m = density*area*h/420*np.array([[156,22*h,54,-13*h],[22*h,4*h*h,13*h,-3*h*h],
                                     [54,13*h,156,-22*h],[-13*h,-3*h*h,-22*h,4*h*h]])
    size = 2*(elements+1)
    mass, stiffness = np.zeros((size,size)), np.zeros((size,size))
    for element in range(elements):
        dofs = np.arange(2*element,2*element+4)
        mass[np.ix_(dofs,dofs)] += m
        stiffness[np.ix_(dofs,dofs)] += k
    values, modes = SparseSymmetricEigenSolver()(mass[2:,2:], stiffness[2:,2:],
                                                 mode_count=3, target=0.0)
    assert np.allclose(modes.T @ mass[2:,2:] @ modes, np.eye(3), atol=1e-8)
    return np.sqrt(values)/(2*np.pi)

beta = np.array([1.875104069,4.694091133,7.854757438])
exact = beta**2*np.sqrt(E*inertia/(density*area))/(2*np.pi*length**2)
counts = [2, 4, 8, 16]
frequencies = np.array([beam_frequencies(count) for count in counts])
print("2 elements [Hz]", np.round(frequencies[0], 4))
print("16 elements [Hz]", np.round(frequencies[-1], 4))
print("analytical [Hz]", np.round(exact, 4))
```

Reference run

```
2 elements [Hz] [  8.386   52.9742 179.1685]
16 elements [Hz] [  8.3819  52.5288 147.0871]
analytical [Hz] [  8.3819  52.5285 147.0813]
```



Convergence of the first three cantilever frequencies under mesh refinement

Download editable notebook · Download Python script

Explain the disagreement between modes

With two elements the first mode is 8.3860 Hz against 8.3819 Hz analytically, an error of about 0.05%. The third mode is 179.1685 Hz against 147.0813 Hz, about 22% high. A conforming displacement space restricts admissible shapes and typically overestimates Ritz frequencies when integration is exact. Shorter spatial wavelengths require more independent shape degrees of freedom. The sixth significant digit of the first mode cannot certify the third.

Why attaching a sensor at a node changes the prediction

An accelerometer of mass m_s at coordinate x_s adds $m_s N(x_s)^T N(x_s)$ to the element mass. It shifts a mode strongly where that mode moves, while its first-order effect at a true modal node vanishes. A location that resolves mode one may be almost blind to mode two.

Exercises on model assembly

Suppress the root deflection but release its rotation. Which diagnostic should fail?

▼ **Worked solution**

A beam with one pinned end and a free end can rotate rigidly with $w = ax$. This mode has zero curvature and stiffness. The smallest eigenvalue should tend to zero. The remaining list must not be compared in order with cantilever roots.

If the beam thickness doubles while width and length stay fixed, predict the frequency change without remeshing.

▼ **Worked solution**

Inertia scales as thickness cubed and area as thickness, so $\sqrt{I/A}$ doubles. All Euler–Bernoulli frequencies double within the slender-beam regime. A thick-beam test can depart systematically owing to shear and rotary inertia.

Test whether a calculated eigenvector is numerically valid without a known exact frequency.

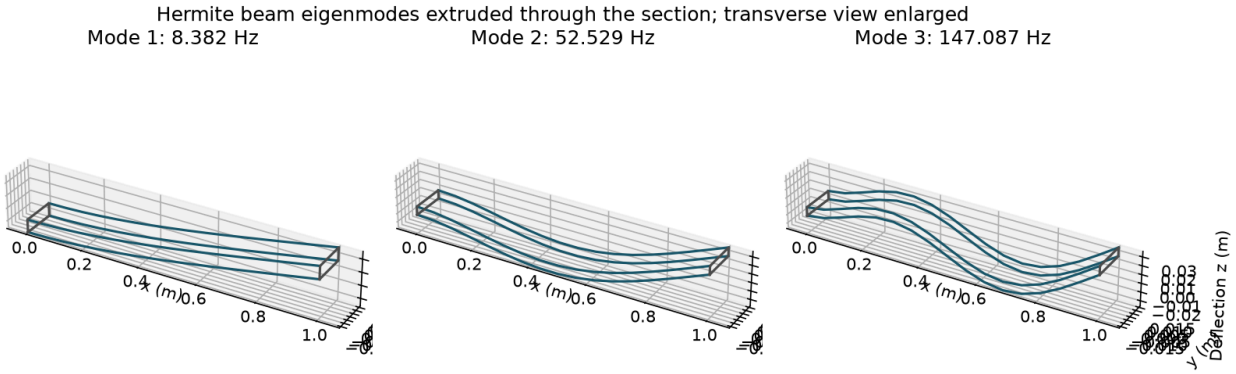
▼ **Worked solution**

Form $K\phi - \omega^2 M\phi$ and scale its norm by $\|K\phi\| + \omega^2\|M\phi\|$. A small residual certifies the discrete eigenpair. Only refinement or physical data test whether the discrete system is adequate.

Further reading · Rao, Mechanical Vibrations, 6th ed., Chapters 7, 8 and 12.

Spatial notebook output · inspect the same solved model in3D

The downloadable notebook now also draws the solved member forces or discrete beam eigenvectors in three dimensions. Its spatial figure below is produced by the actual OneSimulation run.



Recorded spatial output; the beam view enlarges transverse dimensions to make its mode shape legible.

Download the executed notebook with3D output

Vibration and identification in machinery · 7. Which approximation shifted the resonance?

A simulation of a pulsed cantilever contains three separate approximations. Measure the effect of spatial interpolation, retaining finitely many modes and representing a changing force on a time grid.

Project the driven system onto measured modes

Use mass-normalized modes $\Phi^T M \Phi = I$ from the preceding unit. Let a force history $p(t)$ act through a fixed load vector b and measure $y = c^T q$. If modal damping ratios ζ_j are an adequate damping model, projection with $q \simeq \Phi_r z$ gives independent second-order oscillators. Their measured frequency response is

$$\frac{Y(i\Omega)}{P(i\Omega)} = \sum_{j=1}^{r} \frac{(c^T \phi_j)(\phi_j^T b)}{\omega_j^2 - \Omega^2 + 2i\zeta_j \omega_j \Omega}.$$

A mode appears only if the force excites it and the instrument sees it. In a collocated force/displacement pair the numerator is a square. At zero frequency the omitted contributions are then nonnegative and sum to a missing compliance. Test this with a static solve $c^T K^{-1} b$. At a noncollocated sensor the residues can have either sign, especially near an antiresonance.

Treat excitation bandwidth as part of the experiment

A rectangular pulse of duration T has a transform proportional to $T\,\mathrm{sinc}(\Omega T/2)$ with $\mathrm{sinc}(u) = \sin(u)/u$. Its first zero is at $f = 1/T$, though higher-frequency lobes remain. A 5 ms pulse therefore has its main lobe up to 200 Hz, encompassing all three modes of the previous example. A first-mode fit can match the ring-down long after impact while missing the initial peak acceleration. Smoothing the pulse rise changes the high-frequency content of the physical loading.

Distinguish stability from phase fidelity

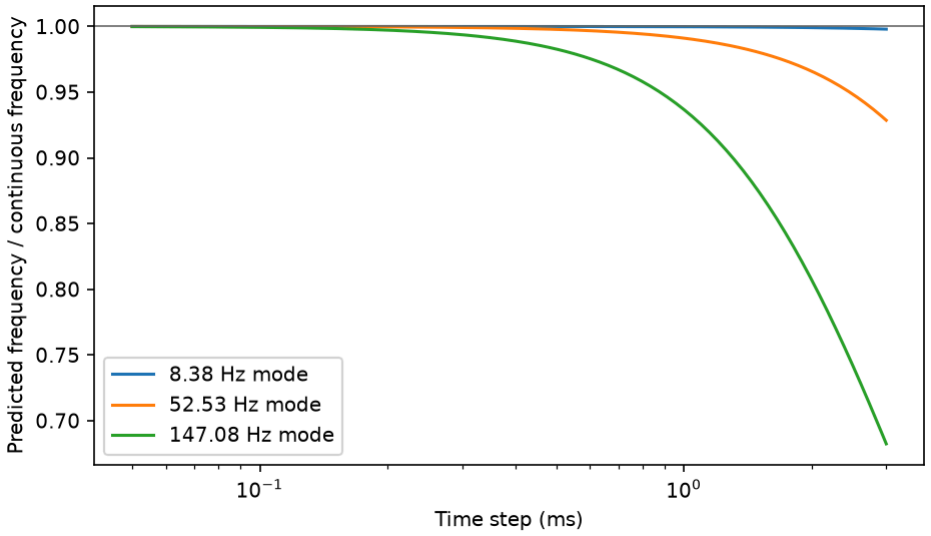
For an undamped mode $\ddot{z} + \omega^2 z = 0$, trapezoidal integration of the two-component first-order system maps its continuous eigenvalue $i\omega$ to the multiplier $(1 + i\omega h/2)/(1 - i\omega h/2)$. Its modulus is one for every step h, but the numerical frequency is

$$\widetilde{\omega} = \frac{2}{h}\arctan(\omega h/2) = \omega - \frac{\omega^3 h^2}{12} + O(h^4).$$

For mode three at 147.08 Hz and $h = 1$ ms, $\widetilde{f}$ is about 137.8 Hz. The signal can remain bounded indefinitely and still be badly out of phase. Halving the step brings it to about 144.5 Hz. At fixed h, refining the mesh alone cannot remove this phase lag. Conversely, integrating the two-element third frequency with an extremely fine time grid merely reproduces the wrong spatial model more faithfully.

How to choose a diagnostic sequence

First converge spatial eigenfrequencies using a modal solver. Next compare the static compliance and measured frequency response when another mode is retained. Finally halve the time step with fixed spatial and modal spaces and a fixed, explicitly interpolated force signal. This keeps the three modeling decisions identifiable.



Trapezoidal time stepping preserves the amplitude of each free mode but retards the phase, particularly for higher frequencies.

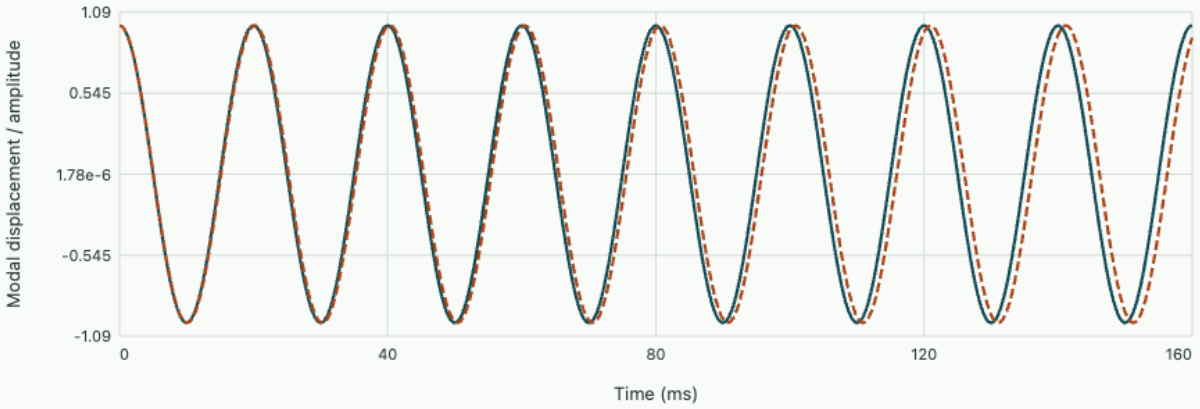
Interactive calculation · Resolve phase error over repeated oscillations

Increase time step at a fixed modal frequency, then compare a higher mode. Trapezoidal integration retains amplitude for this free undamped oscillator while slowing its phase.

Open this live calculation

Try the calculation

Mode frequency (Hz)	50
Time step (ms)	1
Observe for cycles	8



Exact free mode Trapezoidal phase

Numerical frequency 49.595 Hz versus 50 Hz. Relative frequency error -0.811%. The energy-conserving amplitude does not certify phase accuracy.

Exercises on model discrimination

Two modes have frequencies 10 and 40 Hz and equal collocated modal residues. How much of their combined static compliance disappears when the second is omitted?

▼ **Worked solution**

The ratio of second to first is $(10/40)^2 = 1/16$. The second contributes $1/17 = 5.88\%$ of the combined compliance, but it can dominate a narrowband force near 40 Hz.

A colleague matches the measured first frequency by adjusting Young's modulus. Can the two-element third-frequency error also be fixed this way?

▼ **Worked solution**

All frequencies scale with $\sqrt{E}$. Matching the first applies one common scale near one. The third-to-first ratio remains wrong. Use spatial refinement or revise the beam assumptions before interpreting a fitted modulus.

Further reading · Rao, Mechanical Vibrations, 6th ed., Chapters 6–8 and 11–12.