Numerical field laboratory

Turn a prediction into a computational experiment. Derive a reference, solve the field problem, challenge the model with a diagnostic and interpret the residual.

Contents

1. **A converged field solver can still give the wrong voltage**
   Verify a two-dimensional electrostatic calculation with two independent charge patterns. Use symmetry to distinguish a transposed output, a coarse mesh and a poorly converged linear solve.
2. **Measure the capacitance of a continuously graded dielectric**
   Find the exact solution of a variable-coefficient field equation, then estimate the same capacitance from charge, voltage drop and electrostatic energy.
3. **Is a computed screening layer resolved?**
   Use a measured decay length to select a mesh, derive the exact discrete penetration constant and decide whether a discrepancy belongs to the physical screening law or the numerical grid.
4. **What can a mesh measure about a point charge?**
   Distinguish conserved charge and measurable exterior fields from a numerical peak that diverges as the source becomes more localised.
5. **Measure the phase speed of a numerical electromagnetic wave**
   Start with a lossless Maxwell wave, build a second-order time update and compare its actual travel time with what the continuum wave equation predicts.
6. **Move the reference plane of a wave measurement**
   Recover a load reflection coefficient from a noncontact voltage measurement and explain why identical amplitude data can represent physically different terminations.
7. **When a film is thinner than a skin depth**
   Solve a finite conducting layer in complex notation and independently account for its losses using a surface impedance and a volume integral.
8. **Separate diffraction from a finite observation window**
   Recover the far-field angular spectrum of a smooth aperture, then isolate the errors from truncating and sampling the aperture field.
9. **Estimate a cavity linewidth from a finite, biased record**
   Connect the time and frequency descriptions of one damped mode, then diagnose why a weak detector background systematically biases the inferred quality factor.

10. **A magnetic orbit that spirals outward in a perfect vacuum**
Track a charged particle in a prescribed uniform field, derive a rotation update and separate numerical energy growth from a real magnetic force.

11. **When differentiating a measured dipole invents radiated power**
Infer dipole radiation from a sampled displacement and show quantitatively why position noise can dominate a seemingly well-resolved power estimate.

12. **Choose a frequency that identifies a coil's stored and dissipated energy**
Design a small-signal experiment to estimate resistance and inductance, propagate amplitude and phase uncertainty and decide when a distributed field model is needed.

13. **Where should a numerical electrostatic domain end?**
A finite mesh needs an outer boundary even when the physical field extends to infinity. Solve the artificial-boundary error exactly for a dipole and test whether enlarging the numerical domain or matching the exterior field is more effective.

Numerical field laboratory · 1. A converged field solver can still give the wrong voltage

Verify a two-dimensional electrostatic calculation with two independent charge patterns. Use symmetry to distinguish a transposed output, a coarse mesh and a poorly converged linear solve.

Choose a voltage before choosing a solver

A very long grounded tube has a square cross section $0 \le x, y \le L$, with $L = 1$ m. The contents are uniform along the axis. A model of its transverse charge density should predict what a voltage probe placed inside would measure. Three numerical statements that can sound similar are actually quite different: the discrete equations have a small residual, the grid solution approximates the continuous solution, and the physical source describes the tube. We can test the first two in a laboratory problem where the last is specified exactly.

Start by prescribing a voltage that satisfies all four walls. Introduce $s_m(x) = \sin(m\pi x/L)$ and choose $V_*(x, y) = [s_1(x) + g s_2(x)] s_1(y)$ volts. The parameter g is the signed amplitude of the second spatial pattern relative to the first. A probe at the centre sees exactly 1 V for every g, since $s_2(L/2) = 0$. A calculation checked at that point alone could entirely lose the second pattern and still pass. At $x = L/4$ on the centreline, $V_* = 1/\sqrt{2} + g$; at $3L/4$ it is $1/\sqrt{2} - g$. Their difference is $2g$ volts. First write down all three expected readings for $g = 0.5$.

Manufacture the source by differentiating the field

The physical problem uses Gauss's law and the electrostatic relation $\mathbf{E} = -\nabla V$. For uniform vacuum permittivity they yield

$$-\nabla^2 V = \rho/\epsilon_0.$$

A product $s_m(x) s_1(y)$ returns the factor $(m^2 + 1)\pi^2/L^2$ when acted on by the negative Laplacian. Applying this separately to the two patterns gives

$$\rho/\epsilon_0 = (\pi/L)^2 [2 s_1(x) + 5 g s_2(x)] s_1(y) \ \mathrm{V}.$$

The right side has dimensions V/m² after including the factor L^{-2}; multiplication by permittivity gives C/m³. When g becomes large enough, part of the imposed density changes sign. That is permissible for a manufactured source. The calculation represents a solution of the stated Poisson problem. Judging whether a particular apparatus could maintain that charge requires its own charge-transport model.

Uniqueness gives the reference solution special status. If there were a different voltage for the same walls and charge, their difference w would satisfy Laplace's equation with zero boundary value. Integrate $w(-\nabla^2 w)$ over the square. Integration by parts leaves $\int |\nabla w|^2$, because w is zero on the boundary. The integral must vanish; w is constant and its boundary makes that constant zero. This is why agreement with the manufactured voltage is a test of the solver and its input, rather than a choice between two physical solutions.

Background · How can the same source test a second boundary condition?

A smooth manufactured solution that vanishes on every edge tests a homogeneous Dirichlet solve. A code that incorrectly sets every requested boundary voltage to zero may still pass. Add a harmonic lift to the known field on a square of width L. For a constant permittivity the charge depends on the Laplacian, so the exact source is unchanged.

$$\widetilde{V}(x,y) = V(x,y) + V_0\Big(1 + \frac{x}{L}\Big),$$

$$-\epsilon\nabla^2\widetilde{V} = -\epsilon\nabla^2 V.$$

Prescribe the lifted values on all four sides and solve again using the same charge samples. The difference of the computed solutions should be the linear lift, which a centred second difference represents exactly on a uniform grid. A discrepancy points first to boundary-value handling or to the way known values enter the right-hand side. This comparison isolates that error from the truncation error on the original sine modes.

Solve the discrete problem on paper as a second reference

Let $h = L/n$. The standard negative Laplacian has centre coefficient $4/h^2$ and four neighbouring coefficients $-1/h^2$. Substitute the samples of $s_m(x)s_1(y)$. In the x direction the sum of forward and backward neighbours is twice the centre value times $\cos(m\pi/n)$. Collect the corresponding y term. The resulting positive grid eigenvalue is

$$\lambda_{h,m} = \frac{4}{h^2}\Big[\sin^2\frac{m\pi}{2n} + \sin^2\frac{\pi}{2n}\Big].$$

The continuum eigenvalue is $\lambda_m = (m^2+1)(\pi/L)^2$. If the source is the continuous source sampled at the nodes, the discrete voltage coefficient is multiplied by $\lambda_m/\lambda_{h,m}$. Since $\sin z < z$ for positive z, each factor exceeds one on a resolving grid. The faster spatial oscillation has the larger relative bias. The resulting total pointwise error can still change sign because the spatial modes themselves have signs.

This provides three independent calculations: a continuous closed-form field, an exact solution of the chosen difference equations and a numerical linear solve. Agreement of the latter two isolates assembly, indexing and algebraic errors. Comparing either with the continuous solution measures spatial discretisation error. Tightening an iterative tolerance indefinitely cannot remove a difference between $\lambda_{h,m}$ and λ_m.

Run the full field and compare genuinely different observations

Open the live experiment

Try the calculation

Cells per side **16**

Signed second-mode amplitude **0.5**

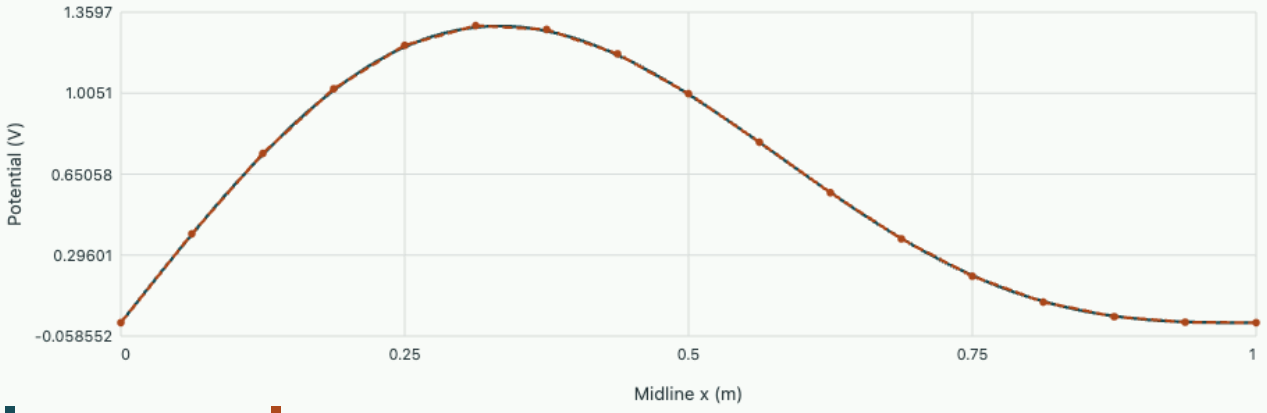

Continuous field Solved grid field

Midline maximum error 0.0077708 V. Quarter-point difference 1.011 V. 2 CG step(s), relative residual 1.45e-13.

Download the OneSimulation model

In your OneSimulation environment, run the model with these inputs.

```
python field-laboratory.py modes --cells 16 --contrast 0.5
```

The browser constructs the right-hand side over the two-dimensional interior and applies conjugate gradients to the five-point operator from an all-zero initial field. The two-mode right-hand side belongs to a two-dimensional invariant subspace, so exact-arithmetic convergence takes at most two iterations. In the special case $g = 0$, only one mode is excited. This unusually small iteration count is a property of the test and says very little about a rough source on the same grid.

Download the model and run the displayed case. Its CSV contains the midline $y = L/2$, ordered by increasing x. For $n = 16,\ g = 2$ the value at $x = 0.1875$ m is about 2.42542 V on the grid and 2.40333 V analytically. If the centre is correct but the 0.1875 m reading is about 0.55736 V, check the orientation of the sampled data array. Keeping x fixed at its midpoint instead of y happens to annihilate the entire second mode.

A convergence report a second person can reproduce

Set $g = 0.5$. Record the centre voltage, the quarter-point voltage difference, the largest nodal error and the relative algebraic residual for $n = 8, 16, 32$. Use even grids divisible by four if reading exact quarter points. Estimate an order from two successive maximum errors using $p = \log(E_n/E_{2n})/\log 2$. For this smooth problem the leading difference between $\sin^2 z$ and z^2 is proportional to z^4, giving a second-

order spatial error at fixed physical modes. If g or the physical source profile changes during refinement, the ratio no longer tests that claim.

An experiment that quotes only the centre voltage has poor sensitivity to g. An experiment quoting only the algebraic residual may perfectly solve the wrong mesh problem. An experiment that compares each plot to itself after smoothing lacks an independent reference. The three diagnostics above address these difficulties with quantities small enough to check by hand.

Extend the verification to a genuinely three-dimensional field

Ground all six faces of a cube with side length L=1 m. A potential that varies in every coordinate is

$$V_*(x, y, z) = A \sin(\pi x/L) \sin(\pi y/L) \sin(\pi z/L),$$

$$\rho = 3\epsilon(\pi/L)^2 V_*.$$

Each second derivative contributes −(π/L)² times the potential. Adding the three derivatives gives the source above; every boundary vanishes because one sine is zero. This is a manufactured source for checking a numerical method. In the supplied calculation ε=1 F/m and A=1 V; scaling both charge and permittivity together leaves the voltage unchanged.

On a uniform nodal grid, the seven-point negative Laplacian has centre coefficient 6/h² and six neighbour coefficients −1/h². The same sine addition identity used in two dimensions gives

$$\lambda_h = \frac{12}{h^2} \sin^2\left(\frac{\pi h}{2L}\right),$$

$$V_h = \frac{3(\pi/L)^2}{\lambda_h} V_*.$$

The browser evaluates this discrete eigenmode directly. The spatial view places a coloured slice inside the cube; its height is the physical z coordinate. Moving that slice reveals the interior field while the companion line plot compares the grid values with the exact three-dimensional solution. The numerical potential is slightly too large because λh is smaller than its continuum limit.

Try it · Rotate a 3D field, move its slice and refine its mesh

Move the plane from the centre to a grounded face and predict the change. Then compare 8 and 16 cells per direction: the maximum nodal error should decrease by approximately four. Doubling A doubles the absolute error while leaving the relative error almost unchanged. Rotation changes the viewpoint only.

Open the live 3D experiment

Try the calculation

Cells per direction	16
Source amplitude (V)	1
Slice position (% of depth)	50

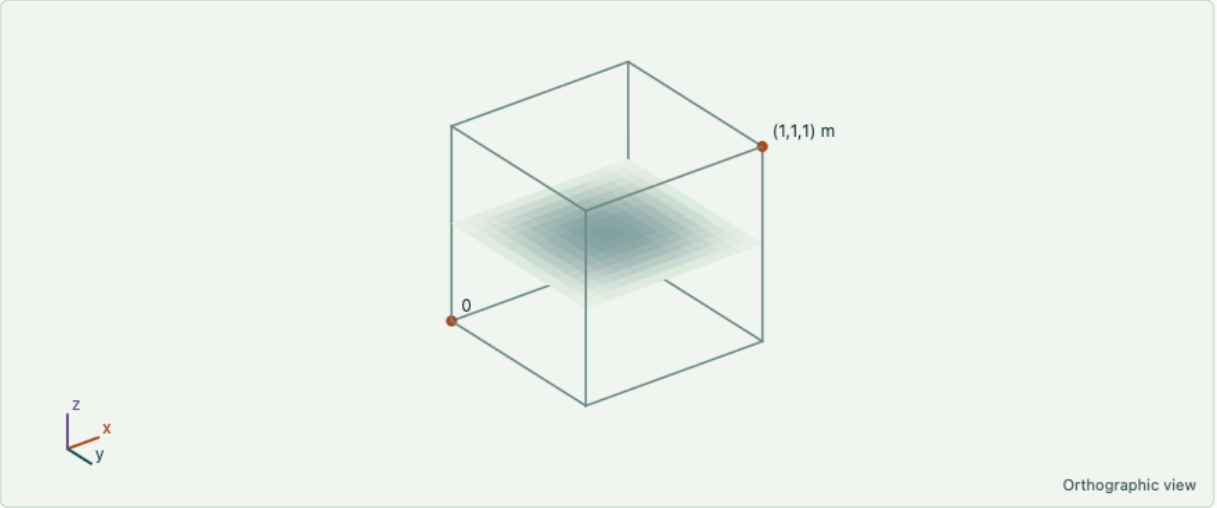

A slice through the three-dimensional potential, at z = 0.500 m. Colour is potential: pale = 0 V, dark green = 2.03 V; slice values range from 0 to 1.003 V. Grid-face colours average their four nodal values. Axes: x = position (m); y = position (m); z = position (m).

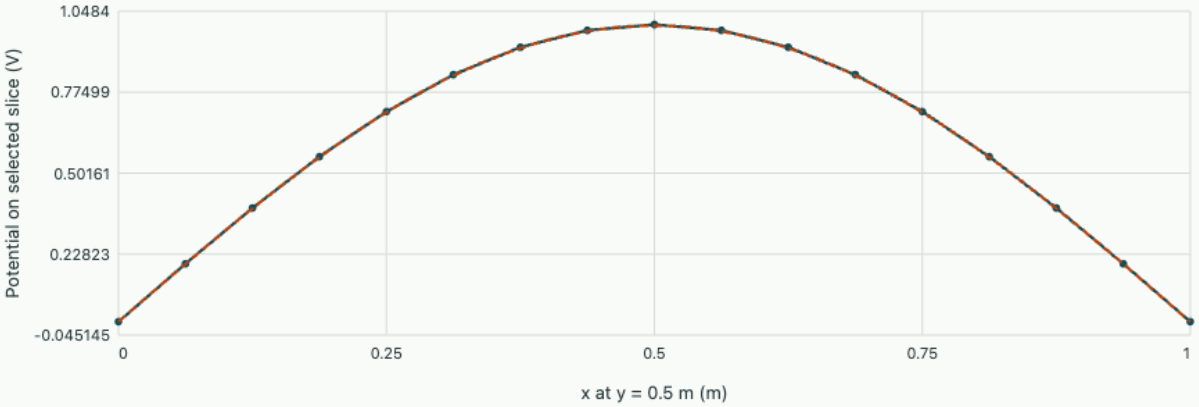

Seven-point FDM · nodal slice Exact 3D solution · same slice

16³ cells; selected grid plane z = 0.500 m. Maximum nodal error in the whole cube 3.219e-3 V (0.322%). All six boundary potentials are zero.

Download the OneSimulation model

In your OneSimulation environment, run the model with these inputs.

```
python poisson-cube-3d.py --cells 16 --amplitude 1 --slice 50
```

Download the 3D OneSimulation notebook · Download Python script

The notebook assembles and solves the full 3D OneSimulation finite-difference model. It reports an RMS error at cell centres; the browser reports a maximum error at nodes for its stated seven-point operator. These are different diagnostics. Compare each against its own continuous reference and repeat with a finer mesh.

```
"""Solve a grounded three-dimensional cube with OneSimulation and verify the field."""
import argparse
from dataclasses import dataclass
import numpy as np
import matplotlib.pyplot as plt
from onesim.cad.d3 import Box
from onesim.meshing import StructuredHexahedralMesher
from onesim.numerics import Numerics
from onesim.numerics.linear import SparseDirect
from onesim.numerics.fdm import FiniteDifference, CentralDifference, CentralLaplacian
from onesim.physics.electrostatics import FixedPotential, PoissonEquation
from onesim.simulation import Problem, Simulation

parser = argparse.ArgumentParser(description=__doc__)
parser.add_argument('--cells', type=int, default=16)
parser.add_argument('--amplitude', type=float, default=1.0, help='Reference potential amplitude in
V')
parser.add_argument('--slice', type=float, default=50, help='Depth as a percentage of the cube')
args = parser.parse_args()
if not 8 <= args.cells <= 24 or not np.isfinite(args.amplitude) or args.amplitude <= 0 or not 0 <=
args.slice <= 100:
    parser.error('Use 8–24 cells, a finite positive amplitude, and a slice between 0 and 100
percent')

@dataclass(frozen=True)
class Permittivity:
    unit = 'F/m'
    def evaluate(self, **state):
        return 1.0

cube = Box(length=1, width=1, height=1)
exact = lambda x, y, z: args.amplitude * np.sin(np.pi*x) * np.sin(np.pi*y) * np.sin(np.pi*z)
field = PoissonEquation(domain=cube.volume, permittivity=Permittivity(),
                        charge_density=lambda x, y, z: 3*np.pi**2*exact(x,y,z))
problem = Problem(name='Three-dimensional manufactured potential', geometry=cube, physics=(field,),
                  boundary_conditions=(FixedPotential(physics=field, boundary=cube.boundaries.all,
value=0.0),))
mesh = StructuredHexahedralMesher(cells=(args.cells,)*3).generate(cube)
simulation = Simulation(problem=problem, mesh=mesh, numerics=Numerics(
    spatial=FiniteDifference(gradient=CentralDifference(), diffusion=CentralLaplacian()),
    linear_solver=SparseDirect()))
simulation.validate().require_valid()
result = simulation.solve_stationary()
potential = result.field(field.potential)
centres = mesh.coordinates[mesh.cells].mean(axis=1)
computed = potential.evaluate(centres)
reference = exact(*centres.T)
error = np.sqrt(np.mean((computed-reference)**2))
z = round(args.cells*args.slice/100)/args.cells
print(f'{args.cells}³ cells; source amplitude {args.amplitude:g} V; slice z={z:.3f} m')
print(f'Cell-centre RMS error against the exact 3D solution: {error:.6e} V')

x = np.linspace(0,1,41)
X,Y = np.meshgrid(x,x)
points = np.column_stack([X.ravel(),Y.ravel(),np.full(X.size,z)])
V = potential.evaluate(points).reshape(X.shape)
```

```
line = np.column_stack([x,np.full_like(x,.5),np.full_like(x,z)])
fig = plt.figure(figsize=(10,4),layout='constrained')
ax = fig.add_subplot(121,projection='3d')
surface = ax.plot_surface(X,Y,V,cmap='viridis',linewidth=0)
ax.set(xlabel='x (m)',ylabel='y (m)',zlabel='Potential (V)',title=f'Computed slice at z={z:g} m')
ax = fig.add_subplot(122)
ax.plot(x,potential.evaluate(line),'o',markevery=2,label='OneSimulation FDM')
ax.plot(x,exact(x,.5,z),'--',label='Exact 3D solution')
ax.set(xlabel='x at y=0.5 m (m)',ylabel='Potential (V)')
ax.grid(alpha=.2); ax.legend()
fig.savefig('poisson-cube-3d.png',dpi=170)
assert np.isfinite(error) and error/args.amplitude < .02
```

Measured output at the default settings:

```
16³ cells; source amplitude 1 V; slice z=0.500 m
Cell-centre RMS error against the exact 3D solution: 1.138076e-03 V
```

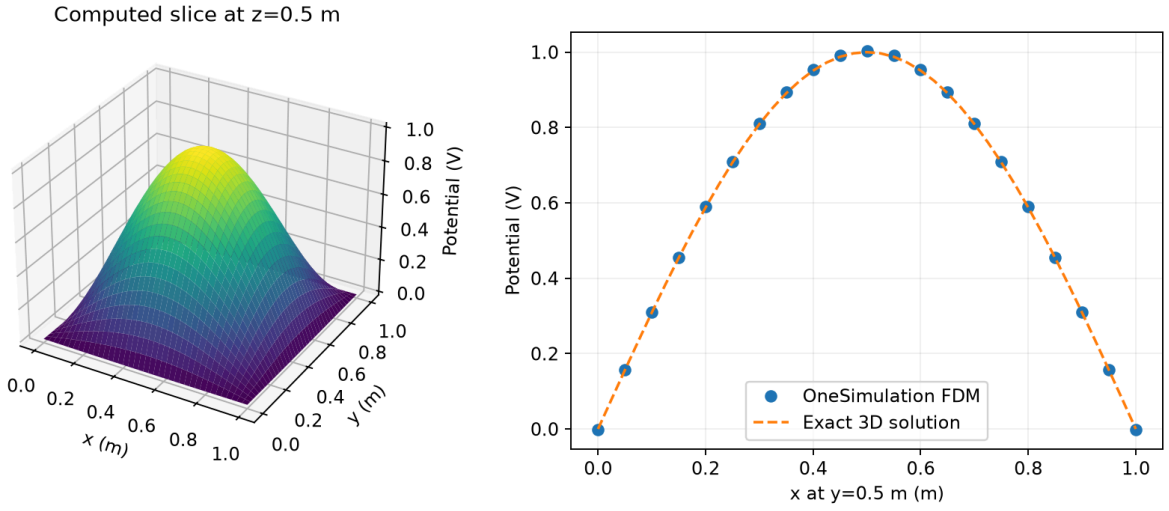

Actual OneSimulation output. Here the left panel plots potential as height above the selected x–y plane; the right panel compares the same field along y=0.5 m.

Tutorial

A colleague sets $g = -0.4$ and reports quarter-point voltages 1.107 V and 0.307 V on a very fine grid. Which change in the recorded setup would explain the observation?

▼ Work through the solution

For the stated sign, the expected pair in increasing x is 0.3071 V then 1.1071 V. Reversing the x axis interchanges the pair while preserving the centre and the symmetric mode. Reversing the sign of the second source mode produces the same pair. A field measurement with a labelled coordinate direction distinguishes these possibilities.

You prescribe the discrete eigenvalue as the source amplitude during every refinement. The returned solution agrees with your analytic curve at all nodes to machine precision. Which part of the verification has disappeared?

▼ **Work through the solution**

This tests the algebra of the discrete eigenmode and the linear solve. It changes the continuum source on every grid and cancels exactly the spatial discretisation error you intended to measure. Rerun with a wavelength and source amplitude specified in physical units independently of the grid.

A second probe has independent voltage noise of standard deviation 5 mV. Estimate the standard deviation of g inferred from the difference of the two quarter-point readings, if both probes have the same independent noise.

▼ **Work through the solution**

The difference has variance $2(0.005)^2$ V² and equals $2g$ volts. Dividing by two gives standard deviation $0.005/\sqrt{2} \simeq 0.00354$ in the dimensionless amplitude. A shared additive offset cancels; a gain error common to both readings does not.

Numerical field laboratory · 2. Measure the capacitance of a continuously graded dielectric

Find the exact solution of a variable-coefficient field equation, then estimate the same capacitance from charge, voltage drop and electrostatic energy.

Build the problem from what is continuous

Two very large parallel plates are separated by $L = 1$ m. Set their potentials to 0 and 1 V. Edge effects are negligible in the observation region. The relative permittivity of the material varies as $\epsilon_r(x) = 1 + gx/L$, where $g \geq 0$. There is no free volume charge. A common first guess is a linear potential between the plates, because the boundary voltages impose an average field. Test that guess using the local field law.

For a thin pillbox whose faces are perpendicular to x, zero enclosed free charge means equal normal displacement at both faces. Let $D_* > 0$ denote the magnitude of the constant displacement. Since potential increases with x, the electric field points towards negative x and $D_x = -D_*$. Hence $V' = D_*/[\epsilon_0(1 + gx/L)]$. Regions of low permittivity must sustain a larger electric field and account for a larger voltage drop per metre. We have a local constraint before solving a single matrix equation.

Obtain the voltage and both expressions for capacitance

Integrate V' from left to right and impose $V(L) - V(0) = \Delta V$. For $g > 0$ this gives

$$\frac{C}{A} = \frac{D_*}{\Delta V} = \frac{\epsilon_0}{L}\frac{g}{\log(1+g)},$$

$$\frac{V(x)}{\Delta V} = \frac{\log(1+gx/L)}{\log(1+g)}.$$

As g approaches zero, both expressions have removable singularities. Since $\log(1+g) = g - g^2/2 + O(g^3)$, the effective relative permittivity approaches $1 + g/2 + O(g^2)$. An arithmetic average can therefore look convincing at very small gradients despite having the wrong general derivation. For computation near zero, use a logarithm specifically evaluated as $\log(1+g)$ with a small-argument routine such as `log1p`.

Check the result through stored energy. Each volume element contributes $\epsilon E^2/2$, so

$$U/A = (D_*^2/2)\int_0^L dx/\epsilon(x).$$

Substituting the voltage relation gives

$$U/A = D_*\Delta V/2.$$

The same capacitance follows from $C/A = 2U/(A\Delta V^2)$. This equality requires a linear, lossless dielectric and a reversible charging path. A material whose permittivity changes with field needs the

corresponding integral over the constitutive curve.

Work one case numerically before assembling the grid

At $g = 3$, the material goes from relative permittivity 1 to 4. The effective value is $3/\log 4 \simeq 2.1640$, smaller than the arithmetic mean 2.5. The midplane voltage is $\log 2.5/\log 4 \simeq 0.6610$ V, not 0.5 V. The electric field at the left is four times as large as at the right. A measurement of the midpoint voltage alone is enough to reject a uniform-field model if the plates and material are well characterised.

As a dimensional check, $C/A \simeq 19.16$ pF/m² for $L = 1$ m. For a geometrically similar layer one millimetre thick the same effective permittivity gives 19.16 nF/m². The voltage profile versus x/L stays fixed when g and the plate voltages stay fixed, while the field and free surface charge increase in inverse proportion to separation.

Assemble by balancing fluxes before dividing by a coefficient

Divide the gap into n cells. Across cell i approximate the displacement magnitude by $D_i = \epsilon_{i+1/2}(V_{i+1} - V_i)/h$. Enforcing $D_i = D_{i-1}$ at each internal node gives a tridiagonal linear system. An independent route views each cell as a capacitor with inverse capacitance per area $h/\epsilon_{i+1/2}$. These series contributions add, giving $D_*^{(h)} = \Delta V / \sum h/\epsilon_{i+1/2}$. Cumulative sums of the same drops then recover the entire grid potential. Both routes implement identical physics and the same discrete constitutive assumption.

The browser uses this cumulative series construction. The OneSimulation input instead supplies the spatial coefficient $1 + gx/L$, zero source and the two exterior potential values to the coefficient-form PDE solver. Their agreement checks assembly against a distinct algorithm. In this example the integral of inverse permittivity also tells us what the midpoint discretisation approximates. Since $1/(1 + gx/L)$ is convex, midpoint quadrature underestimates its integral and therefore overestimates capacitance. This predicts the bias before looking at a convergence table.

Open the live experiment

Try the calculation

Cells across gap 16

Permittivity increase g 3

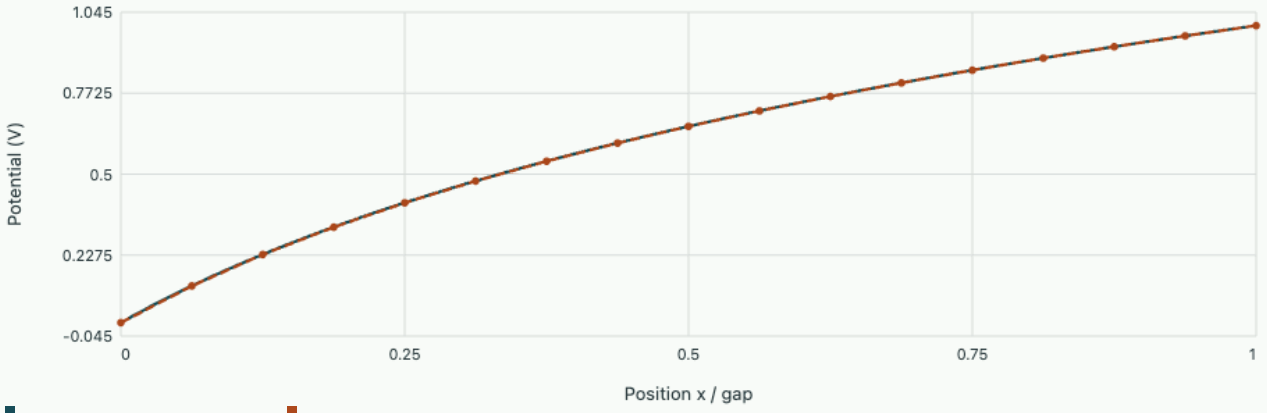

Analytical voltage Flux balance

Effective relative permittivity 2.1662 on the grid and 2.164 analytically. Midpoint voltage 0.66073 V.

Download the OneSimulation model

In your OneSimulation environment, run the model with these inputs.

```
python field-laboratory.py capacitor --cells 16 --contrast 3
```

Compare a local field measurement with an integral quantity

In the experiment, plot V across the dielectric and record the effective relative permittivity. At $g = 2$, $n = 16$, the computed potential at $x = 0.1875$ m is about 0.28974 V; the analytical reference is about 0.28987 V. The potential there is low while the computed capacitance is high. A voltage error at an arbitrary coordinate need not share the sign of the error in an integral quantity.

Now estimate the left electric field from the first-cell drop and from the analytical derivative at $x = 0$. The first-cell drop divided by h is an average across that cell, representing its midpoint to second order for a smooth field. Interpreting it as a value exactly at the wall introduces an $O(h)$ offset. Label sample positions when comparing a derivative against an analytic curve. The grid solution can converge quadratically in voltage while a carelessly placed flux diagnostic appears first order.

Tutorial

The gradient is reversed so $\epsilon_r = 1 + g(1 - x/L)$, but the plates remain at the same voltages. Which earlier observations change?

▼ **Work through the solution**

The integral of inverse permittivity is invariant under $x \mapsto L - x$, so the capacitance and stored energy stay the same. The new potential obeys $V_{\mathrm{rev}}(x) = \Delta V - V(L - x)$. The largest field is now at the right wall, and the midpoint voltage at $g = 3$ falls from about 0.661 to 0.339 V.

A sensor measures an apparently linear voltage profile and C/A greater than the vacuum value. Is the inferred dielectric necessarily uniform?

▼ **Work through the solution**

Under the one-dimensional, source-free and linear assumptions, exactly constant nonzero field and constant normal displacement require constant permittivity. A sparse set of voltage probes may simply undersample a graded profile. Lateral variation, free volume charge or uncertain boundary voltages would break the inference.

Derive bounds on the effective permittivity if only its material range $\epsilon_{\min}$ to $\epsilon_{\max}$ is trusted.

▼ **Work through the solution**

Integrate $1/\epsilon_{\max} \le 1/\epsilon(x) \le 1/\epsilon_{\min}$, multiply by $1/L$, and invert the positive quantities. The effective value lies between these two limits. A simulation reporting an effective permittivity outside that interval in this passive one-dimensional model fails this check regardless of mesh resolution.

Numerical field laboratory · 3. Is a computed screening layer resolved?

Use a measured decay length to select a mesh, derive the exact discrete penetration constant and decide whether a discrepancy belongs to the physical screening law or the numerical grid.

Derive the screened field and state its range

A dilute sample adjacent to an electrode can often be approximated, for small applied potentials, by a relation $\rho(x) = -\epsilon V(x)/\lambda^2$. It says that an excess positive potential attracts net negative charge. The parameter λ is independently calibrated for the ambient material. Combining this relation with Poisson's equation yields

$$V'' = V/\lambda^2.$$

A plus sign is essential: the physical disturbance decays into the sample. Changing the sign would turn a decaying layer into an oscillatory boundary-value problem.

Set $V(0) = V_s$ and $V(L) = 0$. Scale the voltage by V_s, position by L, and write $g = L/\lambda$. The dimensionless problem is

$$-u'' + g^2 u = 0,\ u(0) = 1,\ u(1) = 0.$$

The program solves for unit boundary voltage as a normalisation. For a dilute monovalent electrolyte at room temperature the linearised microscopic model requires the physical $|V_s|$ to be small compared with $k_B T/e$, roughly 26 mV. A computed unit-amplitude solution can represent a small physical voltage by multiplication. Applying 1 physical volt to the linearised electrolyte would generally exceed this assumption.

Solve the finite slab before taking the half-space limit

The general solution contains e^{gx} and e^{-gx}. Impose the right voltage first, using $\sinh[g(1-x)]$, then divide by $\sinh g$ to set the left voltage. Thus $u = \sinh[g(1-x)]/\sinh g$. For $g = 0$ take the limit to get $u = 1 - x$. For a thick slab and fixed physical distance from the left, use $e^{-gx}(1 - e^{-2g(1-x)})/(1 - e^{-2g})$. Close to the right wall the subtracted exponential is needed to satisfy grounding. In the interior many screening lengths away, even a big relative error can correspond to a negligible physical voltage.

At $g = 4$, the midpoint voltage fraction is $1/(2\cosh 2) \simeq 0.1329$. The half-space estimate is

$$e^{-2} = 0.1353.$$

A sample only four screening lengths across is already close to the half-space answer at its middle, though its right edge still has exactly zero voltage rather than e^{-4}. Determine where your instrument can distinguish the profiles before treating one as a better physical model.

Obtain the grid behaviour without running a solver

With grid spacing $h = 1/n$, the interior stencil is

$$-u_{i-1} + (2 + g^2 h^2) u_i - u_{i+1} = 0.$$

Try a geometric progression $u_i = e^{-\gamma_h i h}$. Divide through by the trial value. The sum of its neighbours gives

$$2\cosh(\gamma_h h) = 2 + (gh)^2.$$

Since $\cosh(2z) = 1 + 2\sinh^2 z$, the effective decay constant is

$$\gamma_h = \frac{2}{h}\operatorname{asinh}(gh/2) = g - \frac{g^3 h^2}{24} + O(g^5 h^4).$$

Thus the numerical layer is slightly too broad for a coarse mesh. The relevant resolution is $gh = h/\lambda$ in consistent units, not the cell count alone. Keeping n fixed while increasing g makes a visually unchanged mesh physically coarser. The finite-slab grid reference follows by replacing g with γ_h in the hyperbolic-sine ratio at the nodes.

Separate the physical and computational experiments

Open the live experiment

Try the calculation

Cells across slab **16**

Gap / screening length **4**

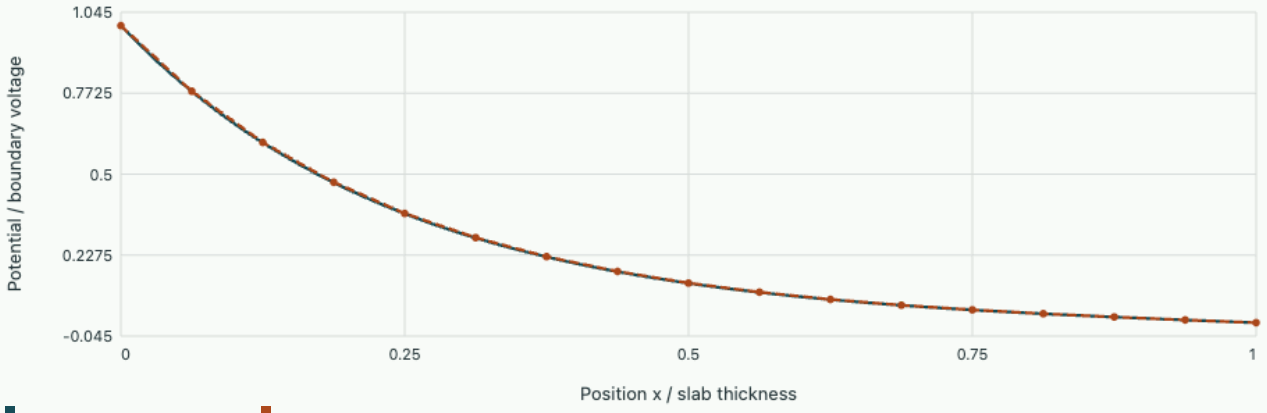

Continuum slab Solved grid values

Midpoint 0.13357, exact 0.1329. Discrete decay constant × thickness = 3.9897.

Download the OneSimulation model

In your OneSimulation environment, run the model with these inputs.

```
python field-laboratory.py screen --cells 16 --contrast 4
```

The browser solves the tridiagonal boundary equations by forward elimination and backward substitution. The downloadable OneSimulation model supplies an absorption term g^2u and the same two Dirichlet values. For $g = 2,\ n = 16$, the potential fraction at $x = 0.1875$ is about 0.67311 on the grid and 0.67297 in the continuum. Reverse the grid spacing in the analytical expression for γ_h and independently predict whether doubling the number of cells raises or lowers the estimate there.

A typical measurement infers a length from the logarithmic slope,

$$1/\lambda_{\mathrm{fit}} = -d\log V/dx.$$

For two nearby positive voltages use the difference of their logarithms divided by probe separation. As V approaches the instrument floor, small absolute noise has a large effect on this ratio. At the grounded right wall the logarithm does not exist. Exclude that point and justify the usable region using both boundary physics and noise. Refining the field mesh cannot improve a voltage measurement below its resolution.

Tutorial

A grid has 16 cells and $g = 16$. Predict the fractional error in the effective decay rate, using the exact discrete expression rather than only the expansion.

▼ **Work through the solution**

Here $gh = 1$, so

$$\gamma_h / g = 2\,\mathrm{asinh}(0.5) \simeq 0.9624.$$

The numerical decay is slow by about 3.76%. The quadratic estimate gives 4.17%, reasonable but visibly imperfect at one cell per screening length.

For a desired relative decay-rate error below 0.1%, estimate the largest acceptable gh in the small-step regime.

▼ **Work through the solution**

The leading relative error is $(gh)^2/24$. Requiring it below 0.001 gives $gh < 0.155$, approximately 6.5 cells per screening length. This condition concerns the exponential rate. The derivative at the electrode may need a separate boundary stencil and its own check.

A calibration reports a longer decay length as electrode voltage increases, even after mesh refinement. Can the numerical dispersion formula alone explain this trend?

▼ **Work through the solution**

In the linear model the normalised profile is independent of the voltage amplitude. The discretisation error depends on g and the mesh. At fixed g, n it cannot create an amplitude-dependent normalised length. Recheck the linear-response regime, electrode boundary condition and instrument response before fitting a new λ.

Numerical field laboratory · 4. What can a mesh measure about a point charge?

Distinguish conserved charge and measurable exterior fields from a numerical peak that diverges as the source becomes more localised.

Ask three detectors different questions

Imagine three instruments looking at a small charged body in vacuum. One measures the voltage on a spherical probe outside it, one measures the integrated electric flux through a closed surface, and one estimates the largest electric field somewhere near the source. If the body is idealised as a point, which readings have finite limits? The distinction matters when a calculation displays a dramatic change in maximum field after a perfectly reasonable mesh refinement.

For a positive point charge q, the exterior potential $V = kq/r$ tends to infinity as the observation point approaches the source. Its radial field magnitude kq/r^2 also diverges. Gauss's law gives a different answer for the flux through any sphere of positive radius:

$$4\pi r^2 E = q/\epsilon_0.$$

A finite integral and an unbounded maximum are compatible. Knowing what the instrument measures determines which numerical output should be expected to converge.

Use a smooth source with a known answer

Replace the potential by $V_a = kq/\sqrt{r^2 + a^2}$, with $a > 0$ a distance. This is a particular physical regularisation. Differentiating before sampling gives

$$E_a = kqr/(r^2 + a^2)^{3/2}.$$

Divide the flux by q/ϵ_0. The putative enclosed-charge fraction becomes

$$F_a(r) = \frac{r^3}{(r^2 + a^2)^{3/2}}.$$

To recover the corresponding density, differentiate $qF_a(r)$ with respect to radius and divide by shell area. The terms involving r^4 cancel, leaving $\rho_a(r) = 3qa^2/[4\pi(r^2 + a^2)^{5/2}]$. This density is smooth, nonnegative and integrates over all space to q, since $F_a(0) = 0$ and $F_a(\infty) = 1$. It has an extended tail. The parameter a is a smoothing scale rather than an outer radius where charge abruptly stops.

At $r = a$ only $1/(2\sqrt{2})$, about 35.36%, of the charge is enclosed. At $r = 2a$ the fraction is about 71.55%. A common intuition that two smoothing distances already contain essentially everything fails for this distribution. Measure the probe radius in units of the source size before blaming the flux calculation.

Hold the observable fixed while refining

The exterior-field ratio relative to the true point charge is $[1 + (a/r)^2]^{-3/2}$. At fixed physical r, expanding the bracket gives $1 - \frac{3}{2}(a/r)^2 + O((a/r)^4)$. Thus a probe far enough away sees a second-order regularisation error. Locating the maximum of E_a, however, requires setting $a^2 - 2r^2 = 0$, giving $r = a/\sqrt{2}$ and a peak proportional to $1/a^2$. Refine a and move the observer with the maximum: the reported peak grows even as a fixed exterior probe reading converges.

Open the live experiment

Try the calculation

Charge-cloud radius a (mm) **3**

Probe radius r (mm) **6**

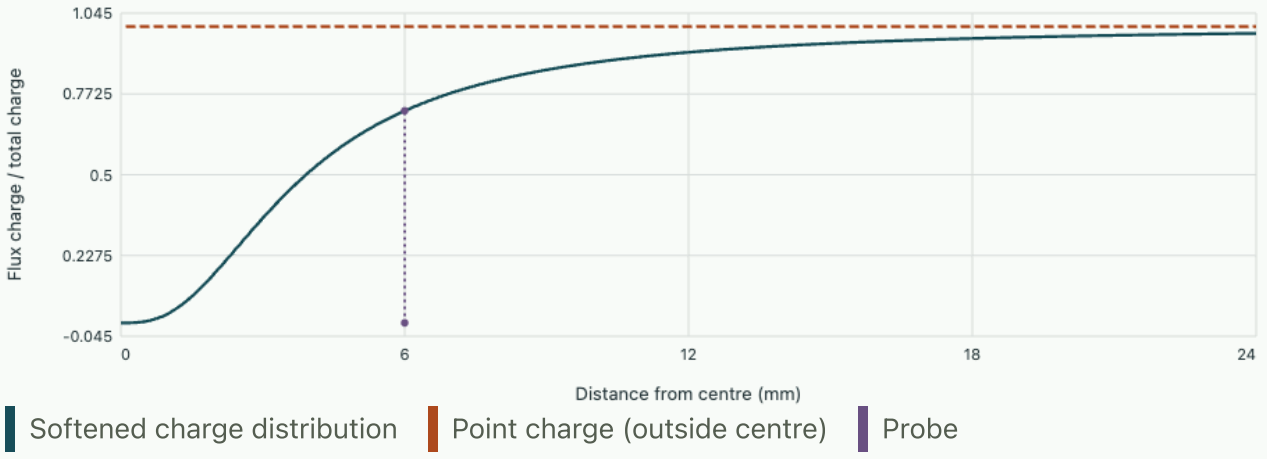

The probe encloses 71.554% of the charge. Reducing a by half gives 91.308% at the same physical probe.

Keep the probe at 6 mm and reduce a from 3 to 1.5 mm. The flux rises from about 71.55% to 91.31% of the point-charge value. Now instead halve both the probe radius and a. The measured fraction stays the same. The two procedures have distinct limiting observations; writing “halve the mesh” without recording the source representation and probe definition would make the experiment irreproducible.

Use continuity as a numerical diagnostic

Suppose a code deposits a prescribed charge into nearby cells and then solves for a nodal potential. Summing deposited cell charges tests normalisation of the source. Summing outward face fluxes over a larger box tests a separate constraint on the computed field. A correct finite-volume implementation cancels interior face pairs. Interpolating a nodal field to a smooth plotting surface with an unrelated quadrature need not conserve the same discrete quantity. A box-flux test using the solver's own numerical face fluxes and a smooth-surface test using the reconstructed field answer complementary questions.

A second check concerns energy. For a true point charge the field-energy integral near the origin scales as $\int_0 r^2 (r^{-2})^2 dr$, which diverges. For a body of radius a this short-distance contribution typically scales like $1/a$. Comparing two source regularisations through their total self-energy without specifying a physical charge radius is therefore an ill-defined validation exercise. Interaction energy with a smooth distant field, in contrast, can have a well-defined point-particle limit.

Tutorial

How far from this smooth charge should a flux probe be placed to enclose 99%?

▼ Work through the solution

Solve $[1 + (a/r)^2]^{-3/2} = 0.99$. This gives

$$r/a = [0.99^{-2/3} - 1]^{-1/2} \simeq 12.2.$$

The answer depends on the chosen source shape. A uniform finite ball would have a different exterior behaviour.

A code uses $a = 2h$ and reports the field one cell away, $r = h$, on every refined mesh. What do you expect for its flux fraction and pointwise field?

▼ Work through the solution

The fixed ratio $a/r = 2$ gives $F = 5^{-3/2} \simeq 0.0894$ at every resolution. The field still scales as r^{-2} times that factor and grows like h^{-2}. The observer is moving towards the singular source.

If total deposited charge differs from q by 2% independently of mesh size, can adjusting a cure the exterior result?

▼ Work through the solution

At large radius the flux follows the total charge that was actually deposited. Changing width cannot repair an incorrect integral. Fix the source normalisation and cell-volume weights before interpreting a convergence slope.

Numerical field laboratory · 5. Measure the phase speed of a numerical electromagnetic wave

Start with a lossless Maxwell wave, build a second-order time update and compare its actual travel time with what the continuum wave equation predicts.

A pulse can arrive late without losing energy

In a uniform lossless one-dimensional material take electric field along y, magnetic field along z and propagation along x. Maxwell's curl laws reduce to

$$\epsilon E_t = -H_x,$$
$$\mu H_t = -E_x.$$

Differentiate the first in time and substitute the second to obtain

$$E_{tt} = c^2 E_{xx},$$
$$c = 1/\sqrt{\mu\epsilon}.$$

A real laboratory signal is a combination of spatial frequencies. Different errors in travel time for different components can distort its shape even when the scheme does not attenuate any resolved Fourier mode.

Use periodic boundaries so that reflection from an end cannot be confused with this effect. Scale distance by one wavelength and time by one exact period. The periodic field $E(x,t) = \sin[2\pi(x-t)]$ has unit amplitude and returns to its initial shape after every integer period. A probe comparing just its envelope would miss a discrepancy in phase.

Derive the discrete dispersion relation

Take points separated by h and times separated by Δt. Centred second differences give

$$u_j^{n+1} = 2u_j^n - u_j^{n-1} + C^2(u_{j+1}^n - 2u_j^n + u_{j-1}^n),$$

with $C = c\Delta t/h$. Insert the complex trial field $e^{i(j\theta - n\Omega)}$. The time pair contributes $2\cos\Omega - 2$, the spatial pair $2\cos\theta - 2$. The physical positive-frequency branch obeys

$$\sin(\Omega/2) = C\sin(\theta/2),$$
$$\frac{c_{\text{phase,grid}}}{c} = \frac{\Omega}{C\theta}.$$

If $0 < C \le 1$ a real Ω exists for every grid phase, and each Fourier mode has constant discrete amplitude in exact arithmetic when started on one travelling branch. For $C > 1$ sufficiently rapid spatial oscillations have no real frequency. One time factor then grows. Testing stability on a single very smooth wave can miss this unstable part of the supported spectrum.

For long waves expand $\sin$ and $\arcsin$:

$$\Omega = C\theta + C(C^2 - 1)\theta^3/24 + O(\theta^5).$$

The relative phase-speed error is $(C^2 - 1)\theta^2/24$ at leading order. Here increasing C towards one improves wave propagation. This special cancellation depends on a uniform one-dimensional grid, constant wave speed and the selected difference operators.

Initialise a clean numerical experiment

A second-order recurrence needs two time levels. Taking $u_j^0 = \sin(j\theta)$ and $u_j^{-1} = \sin(j\theta + \Omega)$ excites a single right-going numerical mode. Using the exact continuum phase $C\theta$ at the second level instead introduces a small mixture of right- and left-travelling grid modes. That is a valid test of initialisation, but it mixes errors if the quantity of interest is solely the numerical dispersion relation. State in your notes how the second level was produced.

Open the live experiment

Try the calculation

Points per wavelength **12**

Courant number **0.8**

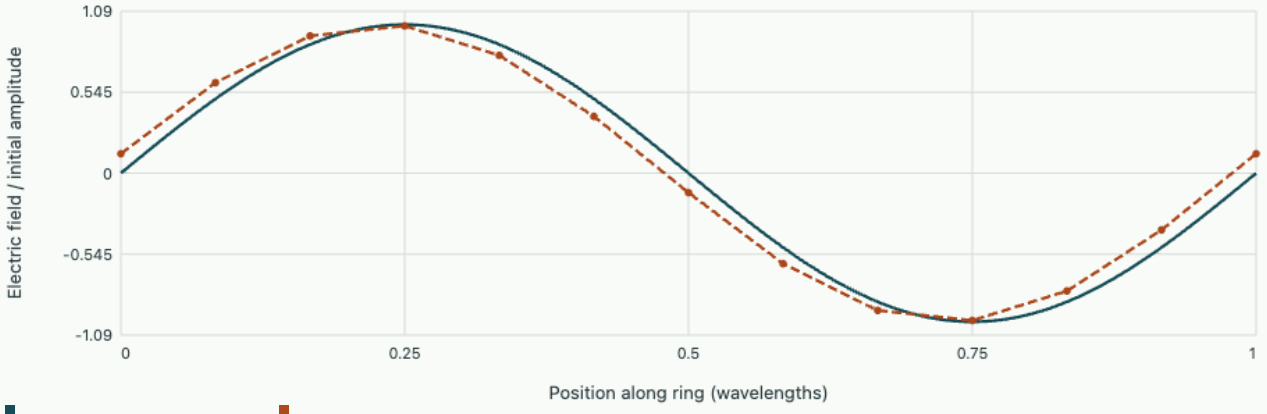

75 steps, 5 physical periods. Phase speed ratio 0.99582. Accumulated phase lag 7.5244°.

The browser actually advances the nodal difference equations. It then reports the dispersion prediction as an independent check. The time is chosen as an integer number of numerical steps close to five exact periods and is stated with the result. At twelve points per wavelength and $C = 0.8$ there are fifteen steps per period. Both exact and discrete solutions have the same amplitude to roundoff. At the end of five periods examine how far their zero crossings have separated.

From a phase test to a pulse test

For a narrow packet, envelope motion is set by group velocity $d\omega/dk$. Differentiate the discrete relation at fixed $h, \Delta t$. The result is $v_g/c = \cos(\theta/2)/\sqrt{1 - C^2 \sin^2(\theta/2)}$ on the positive branch for resolved nonendpoint modes. At $C = 0.8$ a mode close to the grid Nyquist limit has small group speed. A wave carrier and its envelope can show different speed errors. In a real dielectric wavelength dependence may also come from the material itself; a grid-refinement test with unchanged material law separates these contributions.

As an independent benchmark, set $C = 1$. Then $\Omega = \theta$ for $0 \le \theta \le \pi$, and the right-going recurrence translates the samples exactly by one grid cell at each time step. The lab should align both curves even on a surprisingly coarse grid for this particular smooth periodic mode. This checks the sign and indexing of propagation without implying exact accuracy for arbitrary geometry and material interfaces.

Tutorial

A signal must accumulate less than 1 degree of phase error over ten periods. Estimate a required points-per-wavelength count for $C = 0.5$ using the leading error.

▼ Work through the solution

The total continuum phase advance is 20π rad. Requiring $20\pi(1 - C^2)(2\pi/N)^2/24 < \pi/180$ gives $N^2 > 450\pi^2$, or $N \gtrsim 67$. Evaluate the exact dispersion relation to check the chosen integer.

Why is measuring only the field at its crest a poor detector of small phase lag?

▼ Work through the solution

The derivative of a sinusoid vanishes at its crest. A small phase shift changes that value only quadratically; at a zero crossing the change is linear. A phase fit using the entire ring is more robust than either single sample.

The numerical envelope changes after a wall is introduced. What additional tests would separate dispersion from reflection errors?

▼ Work through the solution

First validate travel time on a uniform periodic domain. Next use a single known interface with an analytical reflection coefficient. Move the observation plane and refine the mesh separately. Compare reflected phase and amplitude with an energy-flux balance so a boundary error cannot masquerade as bulk dispersion.

Numerical field laboratory · 6. Move the reference plane of a wave measurement

Recover a load reflection coefficient from a noncontact voltage measurement and explain why identical amplitude data can represent physically different terminations.

Specify which plane defines the load

A uniform lossless transmission line has real characteristic impedance $Z_0 = 50\ \Omega$. Drive it sinusoidally at one frequency. Choose peak phasors with time factor $e^{i\omega t}$ and measure distance d from the load towards the source. Incident and reflected waves move in opposite directions. Taking the load as phase origin, the ratio of the reflected to incident voltage at the load is

$$\Gamma_L = (Z_L - Z_0)/(Z_L + Z_0).$$

The ratio at a sensor plane a distance d away is $\Gamma_{\mathrm{in}} = \Gamma_L e^{-2i\beta d}$. The factor two is the phase difference of two opposite travellers. Moving the sensor does not change a physical resistor.

For a real resistive load, Γ_L is real and ranges from –1 for a short to +1 for an open. A 75 Ω load gives +0.2 and a $100/3$ Ω load gives –0.2. These loads have the same reflected power fraction of four per cent. Distinguishing them requires phase information or positioning the sensor at more than one known location.

Derive the envelope and the ambiguity

The local voltage is the sum of incident and reflected phasors. Factor out the incident phase and take the magnitude:

$$\frac{|V(d)|^2}{|V^+|^2} = 1 + |\Gamma_L|^2 + 2\,\mathrm{Re}\{\Gamma_L e^{-2i\beta d}\}.$$

Its maxima and minima have amplitude ratios $1 + |\Gamma|$ and $1 - |\Gamma|$ for a passive load, giving $\mathrm{VSWR} = (1 + |\Gamma|)/(1 - |\Gamma|)$. The distance between successive maxima is $\lambda/2$, because the measured intensity involves a relative wave number 2β. A measured ratio 1.5 yields $|\Gamma| = 0.2$ but leaves its sign and, for general impedances, its phase undetermined. Extremum positions contain information that a ratio discards.

At one eighth of a wavelength from a 75 Ω load the phase rotation is $-\pi/2$, giving $\Gamma_{\mathrm{in}} = -0.2i$. If a researcher plugs that complex coefficient directly into $Z = Z_0(1 + \Gamma)/(1 - \Gamma)$, the inferred impedance is about $46.15 - 19.23i\ \Omega$. That is the correct input impedance at the sensor plane. To infer the actual element multiply back by $e^{2i\beta d}$ before the inversion. The recovered load is then 75 Ω with zero imaginary part.

Put a probe between voltage extrema

Open the live experiment

Try the calculation

Load resistance (Ω) 75

Probe plane from load (wavelengths) 0.13

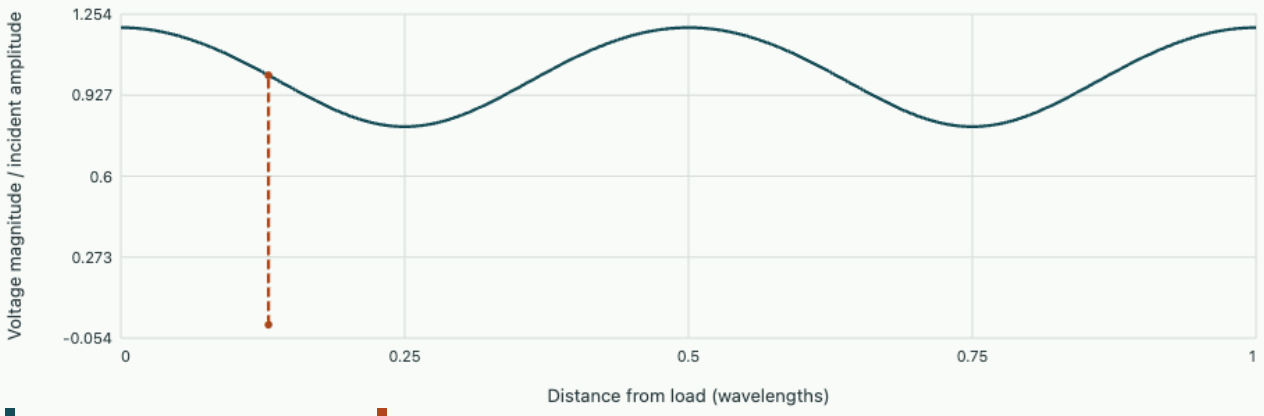

Standing-wave envelope Selected reference plane

Load Γ = 0.2. Probe-plane phase -93.6°. Voltage standing-wave ratio 1.5.

Set the load to 75 Ω and move the sensor from the load through a quarter wavelength. The envelope changes from 1.2 to 0.8 times the incident amplitude. At $\lambda/8$ it is $\sqrt{1+0.2^2}$. Centre your first trial on a voltage extremum and estimate how much the reading changes for a small positioning error. Repeat around the steepest part of the envelope. A high-slope position is sensitive to phase; a stationary amplitude extremum is useful when estimating the reflectance magnitude.

Use the matched case next. The field envelope becomes flat. The program reports that reflection phase is undefined. A zero complex amplitude has no observable phase angle. An instrument that prints a quickly varying phase for a nearly matched load can simply be reporting the angle of noise around the origin. Propagate uncertainty in real and imaginary parts before interpreting such a number.

Recover a complex reflection from several amplitude measurements

If $|V^+|$ is independently calibrated, sample squared voltage at $d = 0, \lambda/8, \lambda/4, 3\lambda/8$. Denote normalised intensity by I_d. Writing $\Gamma = x + iy$ gives $I_0 - I_{\lambda/4} = 4x$ and $I_{\lambda/8} - I_{3\lambda/8} = 4y$ with the convention above. Constant detector offsets cancel in these differences. A multiplicative gain error remains and must be calibrated. A mismatch in the assumed wavelength rotates or distorts the reconstructed coefficient rather than merely adding random noise.

For $\Gamma = 0.2$, the four intensities are 1.44, 1.04, 0.64 and 1.04. The reconstruction returns $x = (1.44 - 0.64)/4 = 0.2$, $y = 0$. The observation probes have provided both quadratures without a phase-sensitive receiver. For a lossy line the round-trip propagation factor also attenuates magnitude; the lossless envelope formula and phase-only de-embedding then require modification.

Tutorial

A reference plane is misplaced by $\lambda/100$. Estimate the resulting phase error of the reconstructed load.

▼ **Work through the solution**

The round-trip phase is $4\pi d/\lambda$. A placement error of $\lambda/100$ therefore rotates the result by 0.04π rad, or 7.2°. Its sign depends on whether distance was over- or underestimated.

Why is recovering a very large resistance from reflection poorly conditioned?

▼ **Work through the solution**

$$dZ/d\Gamma = 2Z_0/(1-\Gamma)^2.$$

Near the open-circuit value +1 this derivative becomes large. A small error in the complex coefficient can imply a large change in estimated resistance while reflected power barely changes.

A standing-wave ratio is measured with a detector displaying power rather than amplitude. Its largest-to-smallest reading is 2.25. What is the correct reflection magnitude?

▼ **Work through the solution**

Take the square root to get a voltage ratio of 1.5, then

$$|\Gamma| = (1.5 - 1)/(1.5 + 1) = 0.2.$$

Substituting 2.25 directly into the voltage formula would yield a biased result.

Numerical field laboratory · 7. When a film is thinner than a skin depth

Solve a finite conducting layer in complex notation and independently account for its losses using a surface impedance and a volume integral.

Which field does the back face support?

A wide conducting film carries an alternating sheet current in the x direction. The field varies only with depth z across thickness t. The surrounding configuration is arranged so the tangential magnetic field at the back face is negligible; at the front its peak phasor is H_s. This back condition describes one particular excitation. A film illuminated from both sides, a nearby magnetic backing or a transmitted wave in a substrate need different conditions. A skin-depth formula by itself does not determine a film impedance.

Use the time convention $e^{i\omega t}$, a real scalar conductivity σ, real permeability μ, and the good-conductor approximation $\sigma \gg \omega\epsilon$. Write $H_y = H(z)$ and $J_x = -H'(z)$. Faraday's law then gives

$$H'' = i\omega\mu\sigma H.$$

Define $\delta = \sqrt{2/(\omega\mu\sigma)}$ and $\gamma = (1+i)/\delta$, so $\gamma^2 = i\omega\mu\sigma$. H and its derivative can be complex even though the observable field is real.

Impose both surfaces before computing power

A solution satisfying the rear condition is

$$H(z) = H_s \sinh[\gamma(t-z)] / \sinh(\gamma t).$$

The electric field at the front is $-H'(0)/\sigma$. Dividing it by the applied field gives the sheet impedance,

$$Z_s = \frac{\gamma}{\sigma}\coth(\gamma t).$$

The thick limit gives $Z_s = (1+i)/(\sigma\delta)$, a resistive and reactive surface contribution of equal size. To find the thin limit expand $\coth w = 1/w + w/3 + O(w^3)$. Substitution gives $Z_s \simeq 1/(\sigma t) + i\omega\mu t/3$. The first term is the ordinary DC resistance per square. The second is the internal inductive reactance for the chosen one-sided magnetic boundary condition. Its factor one third can be checked by integrating the magnetic energy of a linear field profile through the thickness.

Take $\sigma = 4.0 \times 10^7$ S/m and a frequency giving $\delta = 0.2$ mm. The thick-film resistance scale is $1/(\sigma\delta) = 0.125$ mΩ per square. A $20\ \mu\mathrm{m}$ film has $t/\delta = 0.1$; its DC value is 1.25 mΩ per square, already ten times the thick-conductor scale. Treating a very thin film as a half-space can severely underestimate heat production for a specified sheet current.

Derive two independent loss measurements

The inward time-averaged Poynting flux per area equals $\frac{1}{2}\mathrm{Re}(E_s H_s^*) = |H_s|^2 \,\mathrm{Re}\, Z_s/2$ for peak phasors in this geometry. Independently the local cycle-averaged Joule power density is $|H'|^2/(2\sigma)$. Integrate it from front to back. A quick proof of equality avoids pages of hyperbolic algebra: integrate $(H'H^*)' = H''H^* + |H'|^2$ across the film. The term $H''H^* = i\omega\mu\sigma|H|^2$ is purely imaginary. Take real parts, insert $H(t) = 0$, and the boundary term at zero reproduces the surface expression. An unexpected minus sign here often means the outward normal has been used as though it pointed into the sheet.

The imaginary part of the same identity measures reactive magnetic storage. This separation helps diagnose a phase-sign mistake: conjugating every phasor reverses the sign assigned to reactive quantities for the opposite time convention, while average Joule heating stays positive. An output that predicts negative dissipation for a passive film warrants checking definitions before interpreting a material resonance.

Inspect how the back boundary enters a measured field

Open the live experiment

Try the calculation

Thickness / skin depth **1**

Surface current (A/m, peak) **10**

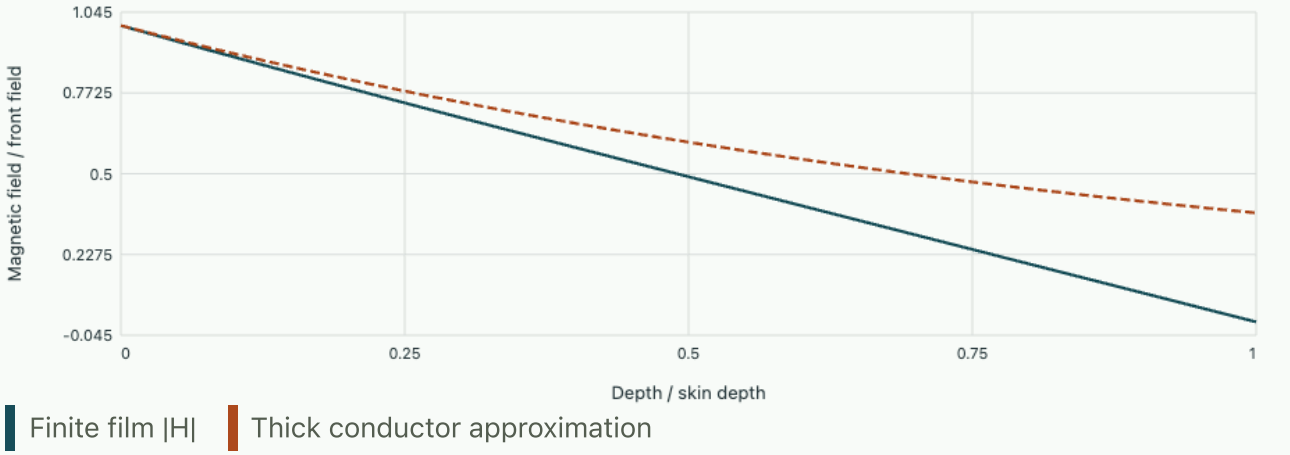

Normalised surface impedance 1.0856 + i 0.65039. Mean power per area divided by the thick-conductor sheet resistance is 54.282 A²/m².

The finite-layer plot uses $|\sinh[(1+i)x]|^2 = \sinh^2 x + \sin^2 x$, with depth in units of δ. The half-space magnitude is simply $e^{-z/\delta}$. At the front both curves are one because the same field is imposed. At the back the finite film must reach zero. Increase t/δ until the front half of the two curves becomes experimentally indistinguishable. Determine whether thickness or the absolute amplitude changes the

normalised profile. Then inspect the reported power, which depends quadratically on the actual surface current.

For $u = t/\delta$, the reported real part in thick-conductor resistance units is $[\sinh(2u) + \sin(2u)]/[\cosh(2u) - \cos(2u)]$. The imaginary part changes the plus to minus in the numerator. Substituting $u = 1$ gives a real part about 1.086. A nominal thickness of one skin depth is quite close to the thick-limit resistance for this boundary problem, even though the magnetic field is forced to zero at the back. Agreement on one integral measurement does not establish that the interior distribution has been measured correctly.

Tutorial

An instrument reports RMS surface fields instead of peak fields. What should change in the Poynting and Joule expressions?

▼ **Work through the solution**

The explicit factor one half is omitted when both phasors are RMS amplitudes. The impedance and the normalised field profile stay unchanged. If an experiment is wrong by a uniform factor two in power, check the amplitude convention before changing material parameters.

At fixed H_s, σ and very small positive frequency, what happens to film loss and to its stored magnetic energy?

▼ **Work through the solution**

In the thin limit the sheet current is nearly uniform, giving mean loss per area $|H_s|^2/(2\sigma t)$ for peak sinusoidal phasors. The peak field decreases approximately linearly from H_s to zero. Its cycle-mean magnetic energy per area is $\mu|H_s|^2 t/12$. The reactive voltage vanishes with frequency although a finite field can store energy.

Could the same Z_s formula be used for an insulating rear surface that imposes zero tangential electric field?

▼ **Work through the solution**

Zero tangential electric field would give $H'(t) = 0$, not $H(t) = 0$. The hyperbolic solution switches to a cosh profile and the sheet impedance uses $\tanh(\gamma t)$. Identify the electromagnetic boundary condition, rather than inferring it just from the word “backing”.

Numerical field laboratory · 8. Separate diffraction from a finite observation window

Recover the far-field angular spectrum of a smooth aperture, then isolate the errors from truncating and sampling the aperture field.

Begin with a transform that can be checked analytically

Consider a long slit system whose transmitted scalar field varies along one transverse coordinate x. At a distant observation point in direction θ, a small aperture element accumulates an approximate relative phase $-kx\sin\theta$. The Fraunhofer amplitude is proportional to $\int A(x)e^{-2\pi i f x}\,dx$ with spatial frequency $f=\sin\theta/\lambda$. The approximation neglects the quadratic phase across the aperture; either a suitable lens or sufficient propagation distance is needed. A numerical Fourier transform cannot cure violation of this optical assumption.

Take the ideal aperture field $A(x)=e^{-x^2/w^2}$ and normalise far-field amplitude by its on-axis value. Integrating by parts or completing the square gives

$$F(f)/F(0)=e^{-(\pi w f)^2}.$$

One derivation avoids any complex contour: differentiate F with respect to f and replace $xA=-(w^2/2)A'$. Integrating by parts gives

$$F'=-2\pi^2 w^2 f F.$$

Solve this first-order equation with the measured $F(0)$. Thus a wider beam on the aperture narrows the far-field distribution. Squaring the amplitude gives the intensity measured by an ordinary detector.

At $w=1$ mm, the $1/e$ amplitude spatial frequency is $1/(\pi w)\simeq 0.3183$ cycles/mm. For 500 nm light this is a small angle of 1.59×10^{-4} radians. An instrument calibrated in radians per millimetre rather than cycles per millimetre would need a factor 2π; stating the Fourier convention prevents that frequent scale discrepancy.

What the instrument has really sampled

A scanner covers $-W\le x\le W$ and gives one amplitude at each of N equally spaced cell midpoints. The numerical transform is the weighted sum $\Delta x\sum_j A(x_j)e^{-2\pi i f x_j}$, where $\Delta x=2W/N$. By symmetry its imaginary part vanishes for a perfectly centred real Gaussian; the program sums the cosine projection. It normalises by the numerical on-axis amplitude, allowing us to compare shapes on a common vertical scale.

Two separate approximations have entered. Truncating the aperture multiplies the field by a rectangle. In frequency space this convolves the Gaussian spectrum with a sinc and can produce signed sidelobes. Sampling at regular intervals makes the spectrum repeat up to a phase convention, with alias copies separated by $1/\Delta x$. Increasing W at fixed N reduces truncation but worsens sampling. Widening the field of view and adding more detector pixels are not interchangeable changes.

Design a convergence test with two axes

Open the live experiment

Try the calculation

Gaussian aperture radius w (mm)	1
Window half-width / w	2
Samples across window	30

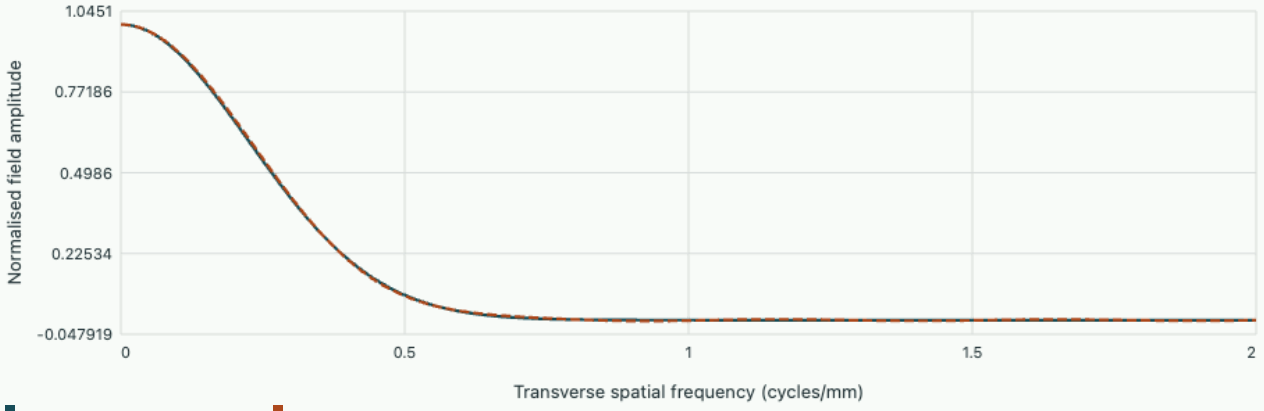

Infinite Gaussian Sampled finite window

Sample spacing 0.13333 mm, Nyquist frequency 3.75 cycles/mm. Measured amplitude at 1 cycle/mm: -0.0014811.

Keep $w = 1$ mm and select 40 samples. Compare $W = w$ and $W = 3w$. The aperture field at the boundary falls from e^{-1} to e^{-9}, so an apparent sidelobe associated with the abrupt window should decrease strongly. Next fix $W = 3w$ and lower the number of samples. The displayed Nyquist frequency falls as Δx grows. When a plotted point lies above it, a smooth-looking value is an aliased sample of the periodic discrete spectrum.

At $W = 2w$, a centred Gaussian retains the on-axis amplitude fraction $\operatorname{erf}(2) \simeq 0.9953$ if integrated exactly. The normalised plot hides this 0.47% loss at zero frequency. If absolute power calibration matters, repeat the experiment without normalisation and independently integrate $|A|^2$ across the aperture. A shape comparison and a throughput comparison are different measurements, both meaningful.

Recover a size from a far-field measurement

Suppose a detector measures the intensity falling to one half at $f_{1/2}$. Since $I/I_0 = e^{-2(\pi w f)^2}$, the inferred w is $\sqrt{\ln 2/2}/(\pi f_{1/2})$. For $f_{1/2} = 0.1874$ cycles/mm the radius is about 1 mm. If the instrument actually reported amplitude half-width the corresponding formula would have $\sqrt{\ln 2}$ in the

numerator; the radius would differ by a factor $\sqrt{2}$. The raw measured quantity and its threshold should accompany any quoted fitted radius.

A positive intensity cannot distinguish a negative sidelobe amplitude from a positive amplitude of the same magnitude. To observe the sign change shown by the finite-window amplitude curve, interfere the field with a phase reference. Conversely, when comparing with a camera image, square the signed numerical amplitude once and apply the detector pixel average. Comparing a real-amplitude plot directly with a recorded intensity array is an avoidable model mismatch.

Tutorial

How could one distinguish an intrinsic physical sidelobe from ringing produced by the scanner window without a known reference aperture?

▼ **Work through the solution**

Increase the scanned width while keeping Δx fixed by adding samples. Window artefacts should change as the omitted tails shrink. Repeat with a different smooth weighting and report its effect on resolution. An intrinsic feature that survives these tests can then be compared against an independent physical model.

At $W = 3$ mm and $N = 24$, what spatial frequencies can be represented without aliasing, and what does adding zero padding do?

▼ **Work through the solution**

$\Delta x = 0.25$ mm, so the Nyquist limit is 2 cycles/mm. Zero padding supplies more displayed frequency samples of the same discrete transform. It adds no measured aperture detail and cannot resolve aliasing already present in the original samples.

A beam shifted by x_0 within a sufficiently wide window has the same far-field intensity. What information changes?

▼ **Work through the solution**

Substituting $x' = x - x_0$ in the Fourier integral multiplies its amplitude by $e^{-2\pi i f x_0}$. The phase ramp records the displacement; the intensity is invariant. A truncated window may break this invariance by clipping different amounts of field.

Numerical field laboratory · 9. Estimate a cavity linewidth from a finite, biased record

Connect the time and frequency descriptions of one damped mode, then diagnose why a weak detector background systematically biases the inferred quality factor.

Know whether the detector reads field or energy

After a drive switches off, let a single cavity mode have complex analytic amplitude $a(t) = a_0 e^{-\kappa t/2} e^{-i\omega_0 t}$ for $t \geq 0$. The energy is proportional to $|a|^2$, giving $U = U_0 e^{-\kappa t}$. Its instantaneous decay rate $-U'/U = \kappa$ implies

$$Q = \omega_0 U / P_{\mathrm{loss}} = \omega_0 / \kappa.$$

A phase-sensitive receiver measuring field amplitude sees a time constant twice as long as the energy lifetime. Before fitting the data, ask how the measurement transducer converts the field into a reported number.

Fourier transforming the causal decaying amplitude gives an integral proportional to $1/[\kappa/2 - i(\omega - \omega_0)]$, up to the time convention. Its squared magnitude has half-power points at $|\omega - \omega_0| = \kappa/2$. Therefore the full power linewidth in hertz is $\Delta f = \kappa/(2\pi)$. A measurement of either the energy decay time $\tau_U = 1/\kappa$ or the spectral width can test the other without having to calibrate the absolute detector gain.

At $f_0 = 2.0$ GHz and $\tau_U = 0.50\ \mu$s, we obtain $\kappa = 2.0 \times 10^6$ s^{-1}, $\Delta f \simeq 318.3$ kHz and $Q \simeq 6283$. A field receiver would report an amplitude time constant of 1.0 µs. Substituting that time directly into $Q = 2\pi f_0 \tau_U$ would double the apparent quality.

Derive the two-point fit with a detector offset

Let a power detector report $S(t) = A e^{-\kappa t} + B$, with a constant background $B > 0$. In the absence of background two samples separated by Δt give

$$\widehat{\kappa} = \log[S(t)/S(t+\Delta t)]/\Delta t.$$

With background, the logarithmic derivative is

$$-\frac{d}{dt}\log S = \kappa \frac{A e^{-\kappa t}}{A e^{-\kappa t} + B}.$$

It starts below the true rate and tends to zero in the tail. A straight line fitted to late log data can therefore look impressively precise while overestimating the lifetime. Increasing the acquisition duration can worsen this bias if the entire record is forced into a straight-line model.

In the experiment set $b = B/A = 0.04$. A two-point fit from one true lifetime to two has dimensionless rate $\log[(e^{-1} + 0.04)/(e^{-2} + 0.04)]$, approximately 0.844. Its apparent lifetime is about 18.5% high. The source, walls and coupling have not changed. Only the interpretation of detector output has

changed. Move the start of the fit towards three lifetimes and inspect the increasing gap from the ideal log line.

Open the live experiment

Try the calculation

Additive detector offset / initial energy **0.04**

Start of fitting window (lifetimes) **1**

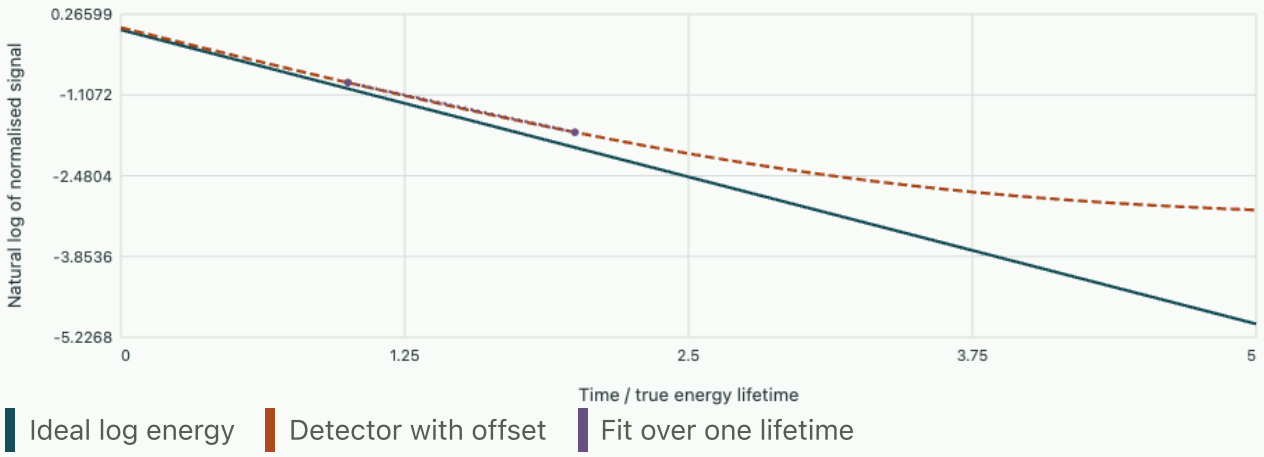

Two-point inferred lifetime 1.1845 times the true lifetime. The same fractional bias transfers to Q for a known mode frequency.

Design a second independent inference

With three equally spaced, noiseless readings S_0, S_1, S_2 of one exponential plus a constant, subtract consecutive values before taking a logarithm. If $a = e^{-\kappa \Delta t}$, the differences are $A'(1-a)$ and $A'a(1-a)$, giving $a = (S_1 - S_2)/(S_0 - S_1)$. This cancels the constant background. At an interval of one lifetime, a first signal of unit energy and background 0.04 gives approximately $1.04, 0.407879, 0.175335$. Their ratio of differences recovers e^{-1}. In a real noise floor two nearly equal readings make this division unreliable. Use multiple windows or fit the raw signal with a measured noise model.

Finite spectral resolution creates a separate challenge. Multiplying the measured field by a rectangular window of duration T convolves its spectrum with a sinc-like response of width of order $1/T$. For the 318 kHz linewidth in the example, $T = 1$ µs gives frequency bins spaced by 1 MHz in a basic Fourier transform. Zero padding can interpolate its curve, but cannot restore the unmeasured late-time signal. A parametric fit may extract a sub-bin frequency or decay rate by using its model assumptions; its uncertainty then depends on signal-to-noise ratio and on whether the one-mode theory is adequate.

Check the single-mode model before refining its fit

If two nearby modes reach the same detector, their fields add before squaring. The result contains an interference term oscillating at their frequency difference on top of a decaying envelope. A scalar positive background cannot explain that alternating residual. Test by narrowing the excitation band, moving the pickup to a symmetry node of one mode, or examining a complex receiver's quadratures. A model with two exponentials has more degrees of freedom and can fit noise as well as signal; require consistency across independent records.

An honest linewidth report specifies the measured quantity, acquisition window, background treatment, coupling state, reference frequency and uncertainty. A fit to stored energy measures total loss. Inferring internal absorption alone additionally requires a known port coupling or a measurement that changes it separately, as in the course's cavity project.

Tutorial

A detector background is independently measured but subtracted too aggressively. What feature can appear in the logarithmic tail?

▼ **Work through the solution**

The corrected signal can approach zero and cross below it. The logarithm then becomes undefined. Even before that point its apparent local slope can become arbitrarily large. Inspect the raw signed residual and its uncertainty instead of dropping nonpositive points without accounting for selection bias.

A second experiment doubles the peak stored energy while the absolute detector offset is fixed. What should happen to the shape of the normalised log tail?

▼ **Work through the solution**

$b = B/A$ halves. Background contamination becomes important later by about $\tau_U \ln 2$. A real change in the physical lifetime need not follow that shift when the response remains linear.

Can a passive resonator's intrinsic losses be determined from a single ring-down measured at one port?

▼ **Work through the solution**

The observed decay is the sum of internal and external rates. A single accurately measured lifetime identifies their sum. A calibrated change in coupling or an independent power budget is needed to separate them.

Numerical field laboratory · 10. A magnetic orbit that spirals outward in a perfect vacuum

Track a charged particle in a prescribed uniform field, derive a rotation update and separate numerical energy growth from a real magnetic force.

Use the work law before solving the trajectory

A nonrelativistic charged particle of mass m and charge q moves perpendicular to a uniform, time-independent magnetic field $B\hat{\mathbf{z}}$. There is no electric field and the dilute test particles do not affect the applied field. The Lorentz force gives $\dot{v}_x = \omega v_y,\ \dot{v}_y = -\omega v_x$ for $\omega = qB/m$. Dotting with $\mathbf{v}$ gives zero, so the speed should stay constant. This property holds for either sign of the charge; the direction of rotation changes. It is a better initial test of a trajectory solver than how circular a short segment looks on screen.

Define $J(v_x, v_y) = (v_y, -v_x)$, so $J^T = -J$ and $J^2 = -I$. The exact velocity over a time step h is $[\cos(\omega h)I + \sin(\omega h)J]\mathbf{v}$. Integration in time gives a circular orbit of radius $v_\perp/|\omega|$ once the guiding centre is chosen. A magnetic field alone bends a trajectory without changing its kinetic energy; a slowly growing orbit radius at fixed B and fixed initial speed demands an explanation.

Why forward Euler adds energy at every step

Euler gives $\mathbf{v}_+ = (I + dJ)\mathbf{v}$ with $d = \omega h$. Squaring the norm eliminates the cross term because $v^T J v = 0$, leaving $\|\mathbf{v}_+\|^2 = (1 + d^2)\|\mathbf{v}\|^2$. No nonzero step has exact unit gain. A geometric picture makes the bias tangible: a tangent step leaves the initial velocity circle on its outside, whatever the point of departure.

After N steps the speed factor is $(1 + d^2)^{N/2}$. With 24 steps per physical orbit, $d = 2\pi/24 \simeq 0.2618$, so a full orbit multiplies speed by roughly 2.2. At shorter times or smaller steps this can resemble weak physical acceleration. Reducing h at fixed physical duration reduces the growth. Testing one step at fixed count, rather than comparing over a fixed physical interval, underestimates the cumulative error.

Solve a centred magnetic step exactly

Use the time-centred velocity in the Lorentz force:

$$\mathbf{v}_+ - \mathbf{v} = dJ(\mathbf{v}_+ + \mathbf{v})/2.$$

Premultiply by the transpose of $\mathbf{v}_+ + \mathbf{v}$. The left side is the difference of squared speeds and the right side vanishes, so the update preserves speed for every d in exact arithmetic. Rearranging yields a two-dimensional linear solve; using $J^2 = -I$ turns it into a rotation,

$$\mathbf{v}_+ = \frac{1 - t^2}{1 + t^2}\mathbf{v} + \frac{2t}{1 + t^2}J\mathbf{v},$$

$$t = d/2.$$

These coefficients are $\cos\alpha$ and $\sin\alpha$ for $\alpha = 2\arctan(d/2)$. This is the magnetic rotation used in the nonrelativistic Boris integrator. For simultaneous electric fields, the usual split step adds two half electric impulses around the magnetic rotation. The current experiment isolates the magnetic part; it makes no claim about a nonuniform electric field or the accuracy of interpolation from a field mesh to particles.

Open the live experiment

Try the calculation

Steps per exact orbit 24

Number of orbits 4

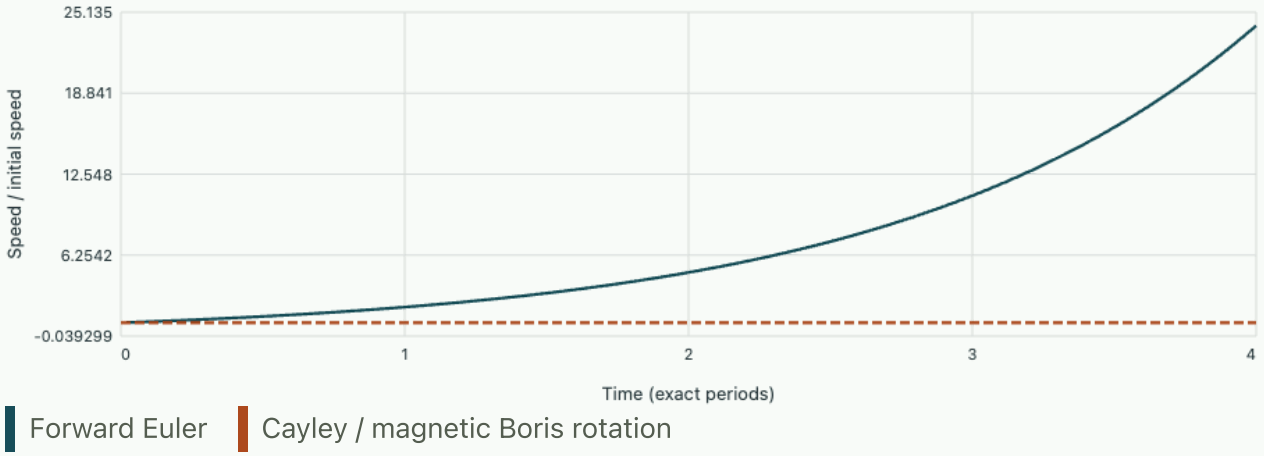

After 4 orbit(s), Euler speed 24.096, rotation speed 1. Rotation still lags the exact phase by 8.1411°.

Speed conservation does not fix arrival time

At small d,

$$\alpha = d - d^3/12 + O(d^5).$$

A speed-preserving orbit lags in gyrophase, even when its energy plot stays perfectly flat. After four true orbits at 24 steps per orbit the predicted lag is about 8.14 degrees. The browser advances the Euler and rotation velocities and reports both norm and lag. The y-axis rescales to fit both curves, so a sharply growing Euler error can flatten the appearance of the well-behaved curve. Inspect numerical values as well as the picture.

Suppose an optical pickup observes the time at which a dilute ion tracer first crosses a known line of sight after several orbits. A phase error can be an important measurement error even when the speed is exact. To use the trajectory as a reference, refine the step until the crossing time changes by less than the experimental timing resolution. If the required integration spans many revolutions, a uniform field

permits exact substeps. In realistic fields one can retain a structure-preserving rotation while resolving spatial variation independently.

A separate diagnostic is time reversal. Starting from the final velocity and applying the rotation with $-d$ exactly inverts a centred step, up to rounding. Forward Euler with $-d$ after $+d$ gives

$$(I - dJ)(I + dJ) = (1 + d^2)I.$$

A forward/backward test detects the lost reversibility without knowing a single experimental constant.

Tutorial

Given an orbit with 48 steps per period, estimate the gyrophase error after ten periods using the leading asymptotic term.

▼ **Work through the solution**

There are 480 steps with $d = 2\pi/48$. The accumulated lag is about $480d^3/12 \simeq 0.0897$ rad, or 5.14 degrees. Doubling the number of steps per period reduces the leading accumulated error fourfold.

A computation preserves speed but reverses the sign of its orbit relative to observation. Would the magnetic work test detect the error?

▼ **Work through the solution**

Both field signs give zero work. Check a signed initial acceleration component $a_y = -qBv_x/m$ for a positive initial v_x, or observe the oriented trajectory. A scalar invariant cannot detect all vector-sign errors.

What changes in the relativistic case with appreciable kinetic energy?

▼ **Work through the solution**

Evolve momentum $\mathbf{p} = \gamma m\mathbf{v}$ and infer velocity from momentum. In a uniform magnetic field without electric work γ stays constant, and the orbital frequency is reduced to $|q|B/(\gamma m)$. An energy-preserving nonrelativistic integration at the wrong frequency would still give the wrong measured period.

Numerical field laboratory · 11. When differentiating a measured dipole invents radiated power

Infer dipole radiation from a sampled displacement and show quantitatively why position noise can dominate a seemingly well-resolved power estimate.

From trajectory to a physically meaningful quantity

An electrically small source carries an oscillating electric dipole moment $p(t) = p_0 \cos \omega t$ in a fixed direction. In nonrelativistic electric dipole radiation, the total instantaneous far-zone power associated with retarded time is

$$P = |\ddot{p}|^2/(6\pi\epsilon_0 c^3).$$

The average over a single sinusoidal cycle is $p_0^2\omega^4/(12\pi\epsilon_0 c^3)$. The fourth power of frequency expresses a physical sensitivity, but a numerical second derivative can introduce a separate sensitivity to high-frequency measurement errors.

A sensor supplies moment or displacement samples at times $t_j = jh$. If the source is a known rigid displacement $x(t)$ of charge q, then $p = qx$. The constant conversion from moment to displacement cancels when comparing a numerical power with a reference power. We therefore scale amplitude by p_0 and time by $1/\omega$. The true acceleration trace is $-\cos t$. Comparing it with the estimated acceleration before squaring reveals signs and phases that a power curve conceals.

Derive the frequency response of the differentiation formula

Use $D_2 p_j = (p_{j-1} - 2p_j + p_{j+1})/h^2$. On a complex exponential $e^{i\nu jh}$ the symbol is $-4\sin^2(\nu h/2)/h^2$. Dividing its magnitude by the exact ν^2 gives a squared sinc factor. At a fixed physical frequency the error is quadratic in h. For a signal oscillating in sign from one sample to the next, $\nu h = \pi$, and the numerical second derivative multiplies its amplitude by $4/h^2$. A numerical derivative smooths the very highest frequency relative to an exact high-frequency derivative yet still strongly amplifies the raw high-frequency component relative to a low-frequency physical acceleration.

Add a deliberately reproducible alternating error $\epsilon(-1)^j$ to the position readings. Its acceleration error is $-4\epsilon(-1)^j/h^2$. Squared and averaged over a complete even grid, the alternating error and the fundamental are orthogonal. Therefore the normalised inferred mean power becomes

$$\frac{\langle P_{\mathrm{est}} \rangle}{\langle P_{\mathrm{true}} \rangle}$$

=\left[\frac{\sin(h/2)}{h/2}\right]^4+2\left(\frac{4\epsilon}{h^2}\right)^2 .

The first term tends to one from below as sampling is refined. At a fixed per-sample alternating displacement amplitude the second grows as h^{-4}. The experiment holds a particular *sampling-related* error fixed; it does not claim that all real noise changes frequency when the sampling rate changes. A fixed-frequency physical disturbance obeys the symbol at its own fixed ν.

A numerical example in which refinement hurts

For 32 samples per period, $h = 2\pi/32 \simeq 0.19635$, and a positional disturbance two parts per thousand of full amplitude gives a spurious acceleration peak about 0.2075 of the true peak. Position data plotted at ordinary scale can look nearly perfect. The normalised mean-square acceleration is already about 1.080. At 64 samples per period with the same alternating error the spurious acceleration amplitude grows fourfold, and its contribution to mean power grows sixteenfold.

Open the live experiment

Try the calculation

Samples per physical period **32**

Position disturbance / amplitude **0.002**

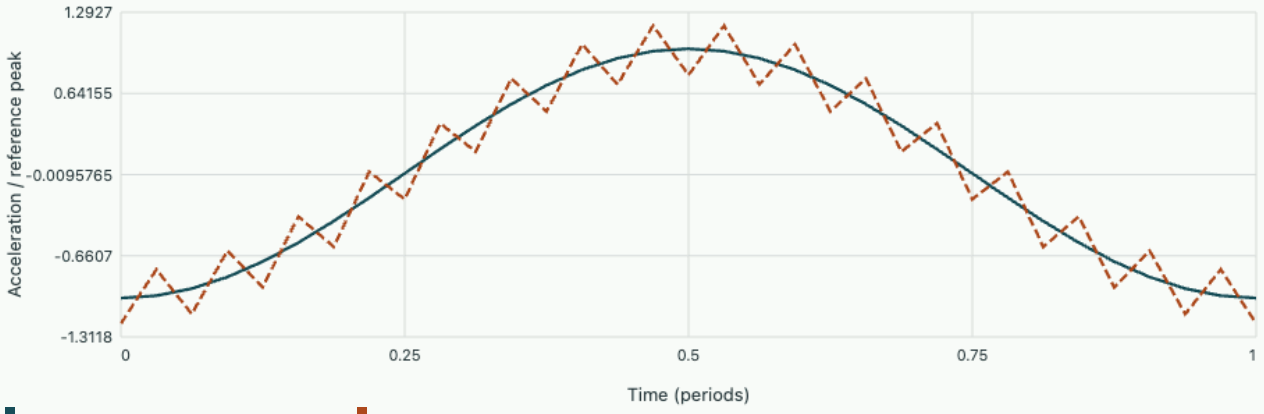

Mean squared acceleration / exact value 1.0797. The alternating error contributes peak acceleration 0.20751 in the same units.

Set the perturbation to zero first and verify the approach to one as the physical sinusoid becomes increasingly well resolved. Restore the perturbation and determine which combination of resolution and disturbance first incurs a ten-per-cent error. The error can initially decrease and later increase with sample rate. Choosing the estimate that happens to lie closest to one at a single resolution can hide a cancellation between physical-mode discretisation error and positive noise power.

Recover a physical model without claiming information you did not measure

If an independent measurement shows the source contains only one sinusoidal frequency, fit its sine and cosine coefficients to the positions rather than differentiating every neighbour. On a full uniformly sampled cycle these coefficients are orthogonal to the alternating offset. Multiplying the fitted phasor by $-\omega^2$ gives the model-based acceleration. This can be a much more precise estimate, but it assumes other harmonics are absent or irrelevant. A weak physical second harmonic has an acceleration

amplification of four and power weighting of sixteen relative to equal moment amplitude. Validate the assumed spectrum before suppressing it.

For broad-band signals one may fit a smooth basis with an explicit high-frequency cutoff or use a model of detector bandwidth. Record the chosen cutoff, its physical justification and the uncertainty propagated into power. Repeating the computation with nearby cutoffs is a test of robustness, not a derivation of the correct filter. Where possible, compare the inferred radiated energy with a distant calibrated flux measurement over the same retarded-time interval. These two measurements have different systematic errors.

Tutorial

Independent, zero-mean random sample errors have variance σ_p^2. What is the variance in one centred second derivative and why need adjacent derivative errors not be independent?

▼ Work through the solution

The weights are $1, -2, 1$ divided by h^2. Independence of the original samples gives variance $6\sigma_p^2/h^4$. Adjacent derivative estimates share two original samples with nonzero weights, so their covariance is generally nonzero.

A second physical harmonic has amplitude 5% of the fundamental. How much does it contribute to the total time-averaged electric-dipole radiation, assuming orthogonal harmonics over the record?

▼ Work through the solution

Its acceleration amplitude relative to the fundamental is

$$0.05 \times 2^2 = 0.2.$$

Its mean power is $0.2^2 = 0.04$ of the fundamental contribution, about 3.85% of their sum. A source spectrum small in displacement need not be equally small in radiated power.

Why must a measured power spectrum specify whether it represents one-sided or two-sided frequencies?

▼ Work through the solution

A real signal contains conjugate spectral amplitudes at opposite frequencies. A one-sided positive-frequency spectrum usually combines their powers. Counting both after applying such a convention doubles the inferred variance and radiation power.

Numerical field laboratory · 12. Choose a frequency that identifies a coil's stored and dissipated energy

Design a small-signal experiment to estimate resistance and inductance, propagate amplitude and phase uncertainty and decide when a distributed field model is needed.

A measured impedance is one complex number

A laboratory sensing coil is driven by a small sinusoidal current at angular frequency ω. In its simplest usable range, a series model predicts $Z = R + i\omega L$. The voltage across a port defines the sign of current entering the positive terminal. With peak phasors, complex power is

$$S = VI^{*}/2 = |I|^{2}(R + i\omega L)/2.$$

Thus the real part measures average irreversible loss. Magnetic energy at peak current is $L|I|^{2}/2$; reactive power equals ω times this peak energy. These equalities supply two routes to the same parameters provided parasitic capacitance and frequency-dependent losses remain negligible.

At 100 Hz, a coil with $R = 20\ \Omega$ and $L = 40$ mH has reactance $2\pi(100)(0.040) = 25.13\ \Omega$, magnitude 32.12 Ω and phase about 51.49 degrees. A 10 mA peak current produces about 0.321 V peak and 1 mW average resistive loss. The peak magnetic energy is only $2\,\mu\mathrm{J}$. Check that the reactive power obtained from the phase and current is about 1.26 mvar. Confusion between RMS and peak phasors easily introduces a factor two while still leaving the impedance ratio correct.

Derive what an amplitude error and a timing error do

Suppose the readout gives magnitude M and phase ϕ. Then $R = M\cos\phi$, $L = M\sin\phi/\omega$. For small independent uncertainties and known ω, differentiate both expressions. The relative first-order perturbations are

$$\frac{\delta R}{R} = \frac{\delta M}{M} - \tan\phi\,\delta\phi,$$

\frac{\delta L}{L}=\frac{\delta M}{M}+\cot\phi\,\delta\phi .

The phase uncertainty must be expressed in radians. For independent random errors their variances add with the squared coefficients; the signed equations above describe individual perturbations. Near zero phase even a very accurate amplitude tells little about inductance, while near 90 degrees inferring resistance requires exceptional phase accuracy. A near-45-degree measurement balances these two sensitivities when the uncertainty budget treats them similarly.

If the phase uncertainty is 0.5° and relative magnitude uncertainty is 1%, the 45-degree standard deviation in either recovered parameter is approximately $\sqrt{0.01^{2} + 0.00873^{2}} = 0.0133$, or 1.33%. At five degrees the relative inductance contribution from phase alone is $\cot 5^{\circ} \times 0.00873 \simeq 10\%$. Taking data at more digits does not improve phase calibration.

Choose the operating point, then challenge the model

Open the live experiment

Try the calculation

Probe frequency (Hz) 100

Coil inductance (mH) 40

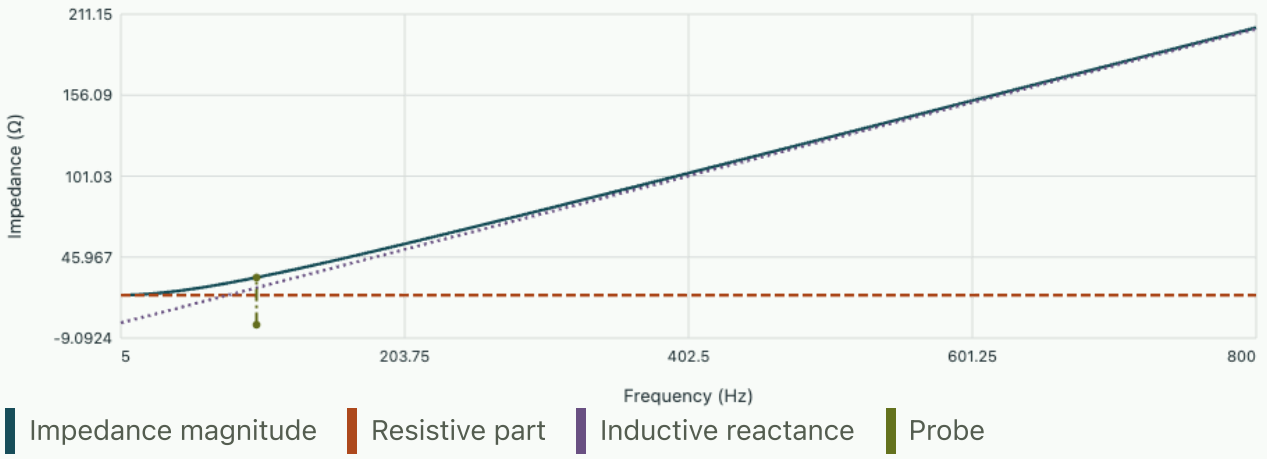

At the probe: |Z| = 32.119 Ω, phase 51.488°. Equal resistive and reactive contributions at 79.577 Hz.

Vary frequency at fixed inductance and read where resistive and reactive contributions are equal. The crossover is $f_c = R/(2\pi L)$, about 79.58 Hz for the default coil. Now double the inductance without changing R; the crossover halves. The graph is a prediction with known parameters. In an actual inverse experiment you know only rough bounds initially. Begin with a small frequency sweep that brackets the crossover, then allocate repeated measurements where the target parameter is most sensitive.

Measure at two or more separated frequencies and reconstruct R and L independently at each. A frequency-dependent apparent R could be caused by skin effect or core loss. A falling apparent inductance near a self-resonance can signal stray capacitance. Simply allowing R, L to take arbitrary values at each frequency fits all data without explaining the mechanism. Add a candidate physical effect and check an independent observable such as loss under a different current amplitude, ring-down frequency or a change in lead geometry.

Close the loop from a fitted parameter to a field calculation

The lumped magnetic energy $LI^2/2$ should match the field integral $\int \mathbf{B} \cdot \mathbf{H}\, dV/2$ for a linear reversible system that includes the relevant surrounding field. This is why the exterior air region and its imposed boundary can affect a numerical estimate even when the measured coil is compact. Increase the domain and refine the mesh separately. A converged mesh in a small box can deliver a very

repeatable value for the wrong boundary problem. The earlier chapter on mutual-inductance matrices generalises this check to cross terms between currents.

Another independent measurement uses a slowly executed current ramp with a known winding resistance. Subtract the measured resistive voltage from terminal voltage and integrate over time. Under the linear model,

$$\int (V - RI)\, dt = L\Delta I.$$

This test weights offsets very differently from sinusoidal phase detection: a constant voltmeter offset integrates into a bias proportional to record duration, while a high-frequency phase error drops out. Agreement of the two methods is stronger evidence than repeating a single instrument setup at the same frequency.

Tutorial

Only voltage and current magnitudes are available, but two frequencies are accurately known. Can a constant R and L be recovered, and what makes the inference fragile?

▼ **Work through the solution**

$$M_i^2 = R^2 + \omega_i^2 L^2.$$

Subtract the equations to get

$$L^2 = (M_2^2 - M_1^2)/(\omega_2^2 - \omega_1^2).$$

Insert this in either equation for R^2 and take positive roots. Nearly equal frequencies produce a small denominator; nearly equal magnitudes amplify measurement noise. Under a misspecified model the inferred square can even become negative.

At a fixed frequency the current reverses direction in a repeat measurement but its reference voltage polarity is held fixed. What measurement error could this create?

▼ **Work through the solution**

Reversing the defined port-current orientation multiplies the inferred impedance by –1 unless the voltage reference is reversed consistently. An apparent negative resistance in passive hardware may be a port convention issue. Check the direction of measured power flow with a known reference resistor.

What extra evidence would justify carrying a fitted inductance into a time-domain prediction?

▼ **Work through the solution**

Make the prediction for a current waveform whose spectrum lies within the validated range of the measurements. Test at another amplitude for nonlinearity or hysteresis and record the residual voltage after subtracting $RI + L\,dI/dt$. A successful fit at one frequency alone does not establish a broadband constitutive law.

Numerical field laboratory · 13. Where should a numerical electrostatic domain end?

A finite mesh needs an outer boundary even when the physical field extends to infinity. Solve the artificial-boundary error exactly for a dipole and test whether enlarging the numerical domain or matching the exterior field is more effective.

Match a harmonic in an empty outer shell

At the surface of a source-containing sphere of radius a prescribe $V(a,\theta) = V_a \cos\theta$. Outside this sphere there is no charge. Insert $V = f(r)\cos\theta$ in Laplace's equation. The angular Laplacian supplies the eigenvalue -2, so $(r^2 f')' - 2f = 0$. A power-law trial gives

$$f = Ar + B/r^2.$$

Decay at infinity removes the growing term and the open-space reference is $f_\infty = V_a(a/r)^2$.

Terminate the calculation at R and ground the artificial wall. The two boundary equations instead give

$$f_R(r) = V_a\Big(\frac{a}{r}\Big)^2 \frac{1-(r/R)^3}{1-(a/R)^3}.$$

A well-converged interior solve can reproduce this answer perfectly while retaining a systematic exterior-domain bias. An alternative boundary condition is $f'(R) + 2f(R)/R = 0$. It admits the decaying dipole and annihilates the unwanted growing coefficient. A single local coefficient matches one spherical-harmonic order exactly.

Estimate the error before refining the mesh

At $r = 2a$ with $R = 5a$, the grounded boundary gives

$$f_R/f_\infty = (1 - 8/125)/(1 - 1/125) = 117/124.$$

The amplitude error is about minus 5.65 percent. At a fixed observation radius, the leading bias decreases with the inverse cube of R. Refining the interior grid at fixed R cannot remove it. Compare a second-order radial difference solution with both analytic curves to measure the two errors independently.

To see what the linear solver actually controls, put a uniform grid on r/a. Approximate the outgoing flux $r^2 f'$ at the right face by $r_{i+1/2}^2(f_{i+1} - f_i)/h$ and form the same quantity at the left face. Subtract the fluxes, divide by h and equate to $2f_i$. The resulting tridiagonal system has positive diagonal $2 + (r_{i+1/2}^2 + r_{i-1/2}^2)/h^2$. Move the known inner and outer boundary terms to the right-hand side. For the grounded wall the outer contribution is zero; the inner value is the amplitude unit.

At $r/a = 2$ the open-domain amplitude is 0.2500. The grounded reference gives 0.2359 when $R/a = 5$. Repeating at $R/a = 10$ raises that prediction to 0.24825. Keep the physical grid spacing

fixed when making this domain comparison. If the cell count stays fixed, increasing the radius also coarsens the mesh, and the two effects can mask each other. The first figure supplies both analytical references needed to separate them.

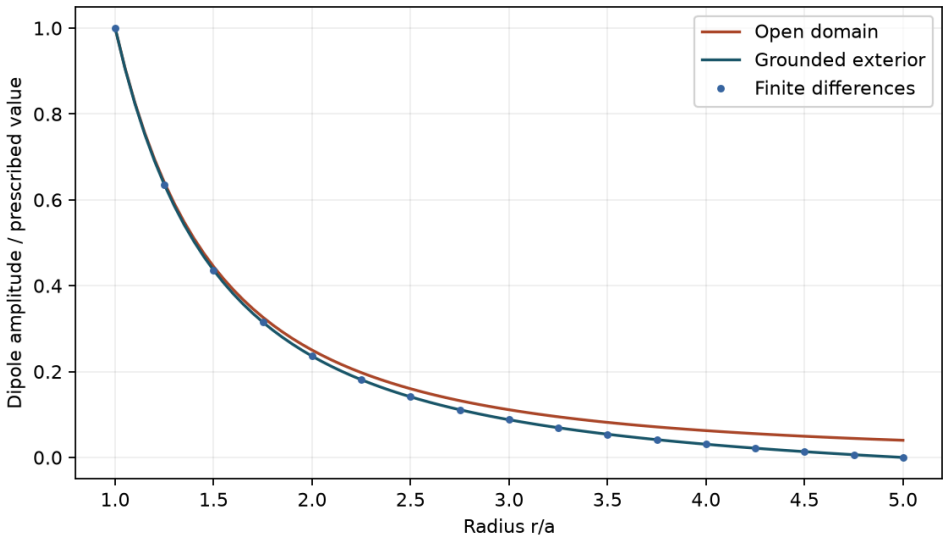

The finite-difference solution follows the field selected by the grounded exterior wall. Its agreement with that curve says nothing about whether the wall models open space.

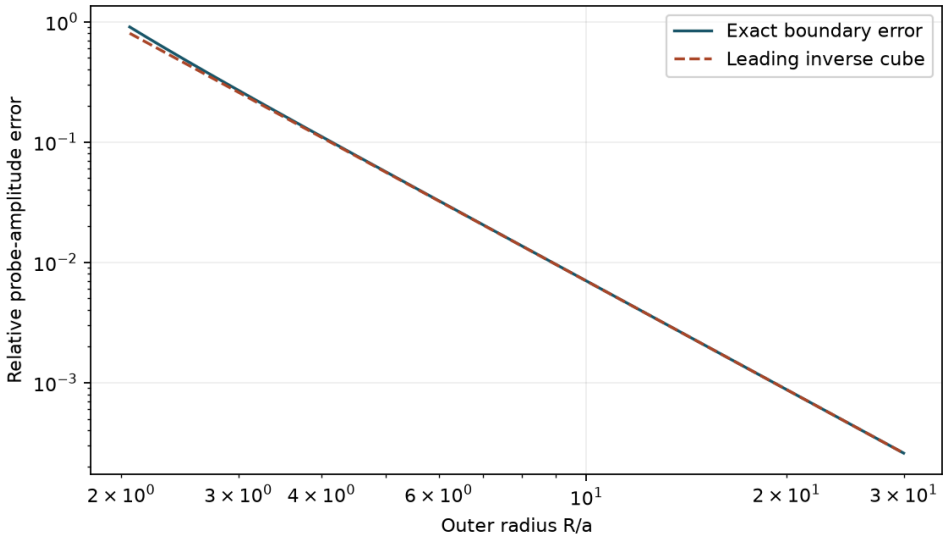

Increasing the domain radius at fixed physical probe location shows the inverse-cube asymptote. A mesh study alone would miss this error.

Exercise · size the exterior shell

For a probe at $2a$, require an error below one percent. What radius is needed?

▼ Work through the solution

The exact relative bias magnitude is $7/[(R/a)^3 - 1]$. Hence $(R/a)^3 > 701$ or $R/a > 8.883$. A value of nine satisfies the condition before numerical discretisation is considered.

Exercise · change the angular dependence

What Robin coefficient would match a pure exterior quadrupole?

▼ **Work through the solution**

A decaying multipole of order ℓ behaves as $r^{-(\ell+1)}$. Its logarithmic slope at R is $-(\ell+1)/R$, so the quadrupole needs $3/R$. A superposition of orders requires angular information or an exterior operator that acts differently on each mode.

Try it · Open domain

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Grounded outer radius / sphere radius 5

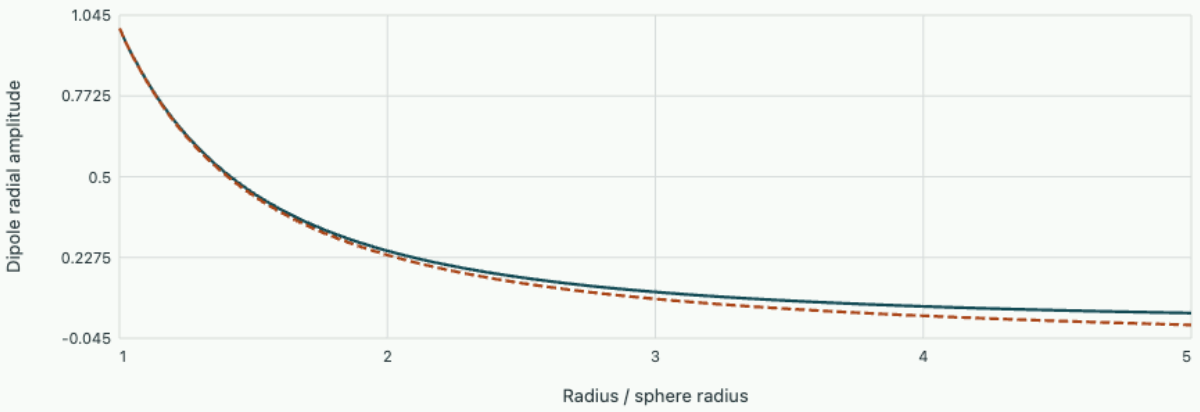

Relative boundary error at r/a=2 is 0.056452. This error remains after solving the discretised bounded-domain problem accurately.

Download the corresponding editable notebook. The browser recalculates the reduced equations directly; the notebook exposes the full calculation and its checks.

Open editable notebook Download .ipynb