

Inverse design of open nanophotonic cavities through differentiable mode expansions

Author list and affiliations in preparation

Working draft — September 5, 2026

Abstract

Modal solvers describe photonic-crystal resonances far more economically than general-purpose field solvers, and differentiating them has opened a fast route to inverse design. The existing differentiable formulations, however, solve a real-frequency eigenproblem on a periodic cell, where the radiative linewidth is not part of the spectrum. We remove that restriction. A finite, one-dimensional nanobeam cavity is written as an outgoing complex-frequency pole of a nonlinear operator $\mathbf{A}(\tilde{\omega}, \boldsymbol{\theta})\mathbf{u} = 0$ assembled from the guided and radiation-continuum spectrum of a reference waveguide and a complete Zernike–Legendre basis inside each etched hole. The outgoing condition is carried by a Sommerfeld contour in longitudinal momentum rather than by an absorbing box, so the imaginary part of the pole—and hence Q —is generated by the continuum itself. Implicit differentiation of the pole through the left and right null vectors returns the complete geometry gradient at the cost of one additional linear solve, independently of the number of geometry variables. We exercise the identical operator protocol on two materially distinct devices: an air-clad silicon Quan–Lončar nanobeam (15 holes per side, 45 variables) and a five-layer $\text{SiO}_2/\text{Si}_3\text{N}_4/\text{anthracene}/\text{PVA}$ cavity with an explicitly subtracted guided feed pole (42 holes per side, 126 variables). Independent finite-difference time-domain (FDTD) evaluation of frozen endpoint geometries gives quality-factor gains of $3.8407\times$ and $3.0666\times$. The central finding is a dissociation: along the second device’s trajectory the surrogate moves from under-predicting FDTD Q (ratio 0.472) to over-predicting it by more than an order of magnitude (ratio 12.31)—a drift of $27.5\times$ in the ratio of the two observables—while the rank ordering of designs within a fixed FDTD protocol is preserved (Spearman 0.911, $n = 23$). Ascent directions transfer; absolute linewidths do not. Both gains are quoted with the exact ringdown conditions under which they were measured, including the endpoint fits obtained over less than one photon lifetime, and are labelled provisional pending a predeclared matched convergence panel.

1 Introduction

Nanophotonic cavity design is an unusually sharp test of an optimization method. The target mode must be confined to a wavelength-scale volume while its weak coupling to a continuum of radiation channels sets a quality factor that can move by orders of magnitude under subwavelength boundary displacements. The quantity being optimized is therefore a small imaginary part riding on a large real one: at $Q = 10^4$ the loss rate is only 5×10^{-5} of the resonant frequency, so a discretization error that is entirely harmless for the resonance wavelength can dominate—and reverse—the Q gradient.

Two families of solver are used for such problems. General-purpose Maxwell solvers optimize arbitrary structures with adjoint methods [1, 2, 3], but every iteration pays for a discretized domain and an absorbing boundary, and the gradient inherits whatever the absorber does to the outgoing field. Modal solvers compress the same physics into a small structured basis; making them differentiable has consequently become an important complementary direction. Guided-mode expansion in particular has been cast in a fully differentiable form and applied to photonic-crystal slab inverse design [4, 5, 6].

That construction solves a real-frequency eigenproblem on a periodic cell. The radiative linewidth is recovered afterwards, perturbatively, from the coupling of a guided mode to the light cone. This is well suited to periodic slabs and is not available for a finite cavity whose entire figure of merit is the linewidth: a closed or truncated modal basis can return a numerically precise and physically meaningless Q , because a handful of guided or absorber modes can reproduce the near field while getting the tiny continuum coupling completely wrong.

The quasinormal-mode (QNM) literature supplies the correct object [7, 8, 9, 10]. A resonance of an open system is a pole of the scattering problem at complex frequency $\tilde{\omega}$ with $\text{Im} \tilde{\omega} < 0$; its normalization is not the closed-cavity energy integral, and substituting the latter into an outgoing problem generally gives the wrong imaginary-frequency derivative even when the real shift looks plausible. What has been missing is a formulation in which the QNM of a *finite, patterned* device is defined by a small nonlinear operator that is (i) built from an analytically known reference spectrum including its continuum, (ii) outgoing

by contour choice rather than by numerical absorption, and (iii) differentiable in closed form with respect to every geometric degree of freedom.

Contributions. This paper reports the following.

- (i) **An outgoing, differentiable mode expansion for finite 1D cavities** (Section 2). The cavity is the root of $\det \mathbf{A}(\tilde{\omega}, \boldsymbol{\theta}) = 0$ for an operator assembled from the exact guided-plus-continuum Green tensor of the unpatterned beam and a complete Zernike–Legendre basis inside each through-hole. Openness is enforced by a Sommerfeld contour in longitudinal momentum, analytically continued to the lower half plane along an explicit homotopy.
- (ii) **A closed-form geometry gradient** (Section 3). Implicit differentiation at a simple pole gives $\partial\tilde{\omega}/\partial\theta_j$ from the left and right null vectors and the analytic derivatives of the hole form factors, so a full 126-variable gradient costs one pole solve plus one reverse-mode contraction rather than 126 forward solves.
- (iii) **Operator-level portability across two devices** (Sections 5 and 6). The same protocol runs on a homogeneous air-clad silicon beam and on a five-layer stack whose reference Green tensor requires Fresnel recursion, Dyson strip resummation, and explicit subtraction of the guided feed pole before the longitudinal contour integral is numerically usable.
- (iv) **A quantified dissociation between direction transfer and magnitude transfer** (Section 7). This is the paper’s most portable result and the one with the clearest consequence for how reduced open-system models should be used.

We are deliberate about what is *not* claimed. The two independently measured Q gains are historical endpoints obtained under stated FDTD settings whose ringdown windows were, in three of four cases, shorter than the fitted photon lifetime (Section 4). They are reported with that condition attached and remain provisional against the predeclared convergence panel of Table 2; no analytical Q is anywhere equated to, or substituted for, an FDTD Q .

2 The cavity as an outgoing pole

2.1 Conventions and the object being computed

We use $\exp(-i\omega t)$ and nonmagnetic, nondispersive media throughout. A passive outgoing resonance is a pole $\tilde{\omega} = \omega_r + i\omega_i$ with $\omega_i < 0$, and

$$Q = -\frac{\operatorname{Re}\tilde{\omega}}{2\operatorname{Im}\tilde{\omega}} = -\frac{\omega_r}{2\omega_i}. \quad (1)$$

The optimization target is $\log Q$, for which

$$\delta \log Q = \frac{\delta\omega_r}{\omega_r} - \frac{\delta\omega_i}{\omega_i}, \quad (2)$$

so that multiplicative gains are additive and the two parts of the pole enter on an equal relative footing. Equation (2) also states the difficulty of the problem plainly: since $|\omega_i|/\omega_r = 1/(2Q)$, the second term is the one that must be resolved, and it is the smaller of the two by four orders of magnitude at $Q \sim 10^4$.

2.2 Exact parameterized dielectric

The unpatterned beam of width W and thickness T defines the reference permittivity $\epsilon_0(y, z)$. Hole j is the elliptical through-cylinder

$$H_j(\boldsymbol{\theta}) = \left\{ (x, y, z) : \frac{(x-x_j)^2}{r_{x,j}^2} + \frac{y^2}{r_{y,j}^2} \leq 1, |z| \leq \frac{T}{2} \right\}, \quad (3)$$

so that $\epsilon_{\boldsymbol{\theta}} = \epsilon_0 + \Delta\epsilon_{\boldsymbol{\theta}}$ with $\Delta\epsilon_{\boldsymbol{\theta}} = (\epsilon_c - \epsilon_b) \sum_j \mathbf{1}_{H_j}$. The geometry vector $\boldsymbol{\theta}$ collects the positive-side centres x_j and the two semi-axes $r_{x,j}, r_{y,j}$ of every hole; the mirror image is generated exactly. This parameterization is not a discretized permittivity map: the Fourier factor of a hole is available in closed form,

$$\hat{\mathbf{1}}_{H_j}(\mathbf{k}) = e^{-ik_x x_j} \left[2\pi r_{x,j} r_{y,j} \frac{J_1(\rho_j)}{\rho_j} \right] \left[\frac{2\sin(k_z T/2)}{k_z} \right], \quad (4)$$

with $\rho_j^2 = (k_x r_{x,j})^2 + (k_y r_{y,j})^2$, and so are its derivatives, e.g. $\partial_{x_j} \hat{\mathbf{1}}_{H_j} = -ik_x \hat{\mathbf{1}}_{H_j}$. Every geometry derivative used later is therefore analytic rather than a finite difference of a remeshed structure.

2.3 Outgoing volume-integral pole equation

With $\mathcal{L}_0(\omega) = \nabla \times \nabla \times - (\omega^2/c^2)\epsilon_0$ and $\mathbf{G}_{\text{wg}}^+(\omega) = \mathcal{L}_0(\omega)^{-1}$ taken on the outgoing sheet, the source-free Maxwell problem is equivalent to the homogeneous Lippmann–Schwinger equation $\mathbf{E} = (\omega^2/c^2) \mathbf{G}_{\text{wg}}^+ \Delta\epsilon_{\boldsymbol{\theta}} \mathbf{E}$. Restricting the unknown to the finite hole support gives the exact pole condition

$$D(\omega, \boldsymbol{\theta}) = \det_{\text{reg}} \left[I - \frac{\omega^2}{c^2} \mathbf{G}_{\text{wg}}^+(\omega) \Delta\epsilon_{\boldsymbol{\theta}} \right] = 0. \quad (5)$$

The regularized determinant is written deliberately. The electric dyadic Green tensor carries a singular local term; an unregularized collocation determinant of the same kernel is not the Fredholm determinant of a trace-class operator and does not have (5) as its root. We keep the longitudinal projector in Fourier space, where it is an algebraic term rather than a $1/R^3$ singularity requiring a depolarization guess (Supplementary Section S4).

2.4 Reference Green tensor: guided poles and continuum

The unpatterned beam is invariant along x , so

$$\mathbf{G}_{\text{wg}}(x - x'; \omega) = \int_{c_\beta} \frac{d\beta}{2\pi} e^{i\beta(x-x')} \tilde{\mathbf{G}}_{\text{wg}}(\beta; \omega). \quad (6)$$

Differentiable outgoing mode-expansion workflow




Figure 1: Operator flow. The reference waveguide supplies an exact guided-plus-continuum Green tensor; the finite hole array supplies a complete interior basis; their contraction on an outgoing longitudinal contour defines $\mathbf{A}(\tilde{\omega}, \boldsymbol{\theta})$, whose root is the cavity pole. The geometry gradient is obtained by implicit differentiation at that root. FDTD is an independent, sparse qualification layer and never supplies an optimizer gradient.

At fixed β the transverse problem has compact support. For a homogeneous cladding the symbol is exact, $\mathbf{G}_c = P_T/(k^2 - k_c^2) - P_L/k_c^2$ with $P_L = \mathbf{k}\mathbf{k}^T/k^2$, and the beam cross-section is restored by a Dyson resummation over its own susceptibility χ_b ,

$$\widetilde{\mathbf{G}}_{\text{wg}}(\beta, \omega) = [I - k_0^2 \chi_b \widetilde{\mathbf{G}}_c(\beta, \omega)]^{-1} \widetilde{\mathbf{G}}_c(\beta, \omega). \quad (7)$$

Guided modes are therefore not inserted by hand: they are the zeros of $\det[I - k_0^2 \chi_b \widetilde{\mathbf{G}}_c]$, i.e. poles of the same resolvent that carries the radiation continuum, and their residues are normalized by the pole derivative rather than by an inconsistent finite-box norm.

For the multilayer device the homogeneous symbol is replaced by the stratified one. Generalized Fresnel reflections are composed recursively through each film, $R_{j-1} = (r_{j-1,j} + R_j e^{2ik_{zj}d_j})/(1 + r_{j-1,j} R_j e^{2ik_{zj}d_j})$, the vertical kernel in the source layer is $g(z, z') = -u_L(z_<)u_R(z_>)/W_r$ with Wronskian $W_r = 2ik_z(e^{-ik_z d} - r_b r_t e^{ik_z d})$, and the patterned strip is added by the same Dyson step as (7). Replacing the films by an effective cladding would delete their guided and leaky poles and is not an admissible reference for Q optimization (Supplementary Section S6).

2.5 Outgoing contour and analytic continuation

Openness enters through $\mathcal{C}_\beta$, not through a numerical box. In the scaled variable $u = \beta/k_c(\omega)$ the radiation branch points are pinned at $u = \pm 1$ and the guided poles sit near $u = \pm n_{\text{eff}}/n_c$. The implementation uses the Sommerfeld path

$$u(t) = t - ih \tanh(t/d) \exp[-(t/L)^8], \quad (8)$$

which passes below every positive- u singularity and above its negative- u partner, and reverts to the real axis outside $|t| \sim L$. Holding $u(t)$ fixed while moving ω from the causal upper half plane down to the resonance is an explicit continuation homotopy; a pole or branch point crossing the frozen path is detected as a discontinuity rather

than silently changing sheet. Substituting a lower-half-plane frequency into a real-axis quadrature lands on the wrong sheet, and did so in an early revision of this work (Supplementary Section S18).

2.6 Interior basis and assembly

Inside hole j we use normalized hole coordinates $x - x_j = r_{x,j} \rho \cos \theta$, $y = r_{y,j} \rho \sin \theta$, and expand in a complete Zernike–Legendre vector basis,

$$\phi_{j,nm\sigma\ell s} = \frac{Z_{nm}^\sigma(\rho, \theta)}{\sqrt{r_{x,j} r_{y,j} N_{nm}}} \sqrt{\frac{2\ell+1}{T}} P_\ell(2z/T) \mathbf{e}_s, \quad (9)$$

with $R_n^m(\rho) = (-1)^{(n-m)/2} \rho^m P_{(n-m)/2}^{(m,0)}(1-2\rho^2)$ evaluated in Jacobi form for stability. Because the hole spans the full beam thickness and shares the waveguide’s Legendre functions in z , the thickness projection is exact. Denoting by $C_j(\beta)$ the projection of this basis into the cross-section basis at longitudinal momentum β , the projected outgoing Green matrix is

$$\mathbf{G}_H^+(\omega) = \int_{\mathcal{C}_+} \frac{d\beta}{2\pi} C(-\beta)^T \widetilde{\mathbf{G}}_{\text{wg}}(\beta, \omega) C(\beta), \quad (10)$$

and the finite operator whose root is the cavity is

$$\mathbf{A}(\omega, \boldsymbol{\theta}) = I - k_0^2(\epsilon_c - \epsilon_b) \mathbf{G}_H^+(\omega, \boldsymbol{\theta}). \quad (11)$$

The left factor in (10) is the transpose $C(-\beta)^T$, not the Hermitian conjugate $C(\beta)^\dagger$. The two coincide on the real-frequency boundary for real basis functions, but only the transpose is analytic in ω ; using the conjugate destroys the continuation into the lower half plane. Reciprocity then predicts a complex-symmetric $\mathbf{A}$, which is a stringent numerical check rather than an assumption (Section 4).

Truncating (9) at radial order N_{xy} and vertical order N_z and the contour at quadrature order N_β yields a sequence $\mathbf{A}^{(N)}$ whose roots $\tilde{\omega}_N$ converge to the cavity pole; N_β is a continuum quadrature order, not a finite number of physical radiation modes.

3 Geometry gradient of the pole

Let $\mathbf{u}$ and $\mathbf{v}$ be right and left null vectors of the assembled operator at an isolated simple pole,

$$\mathbf{A}(\tilde{\omega}, \boldsymbol{\theta}) \mathbf{u} = 0, \quad \mathbf{v}^T \mathbf{A}(\tilde{\omega}, \boldsymbol{\theta}) = 0. \quad (12)$$

Differentiating (12) with respect to a geometry coordinate θ_j and contracting on the left with $\mathbf{v}$ eliminates the unknown field derivative, giving

$$\boxed{\frac{\partial \tilde{\omega}}{\partial \theta_j} = - \frac{\mathbf{v}^T (\partial_{\theta_j} \mathbf{A}) \mathbf{u}}{\mathbf{v}^T (\partial_{\omega} \mathbf{A}) \mathbf{u}}} \quad (13)$$

The denominator is the QNM normalization for this operator. It is generated by the frequency dependence of $\mathbf{A}$ itself—including that of the reference Green tensor and of the contour—and any arbitrary rescaling of either null vector cancels between numerator and denominator. This is the safe interface: a closed-cavity energy normalization substituted into an outgoing problem generally returns a wrong $\partial \omega_i / \partial \theta_j$ even when $\partial \omega_r / \partial \theta_j$ looks correct, and it is precisely ω_i that carries Q through (2).

Three properties make (13) practical.

Cost. The denominator and both null vectors are computed once per accepted geometry. Each additional coordinate costs one contraction of the already-assembled $\partial_{\theta_j} \mathbf{A}$, whose ingredients are the analytic form-factor derivatives of (4). Reverse-mode accumulation over the operator primitives therefore returns the entire gradient—45 coordinates for the first device, 126 for the second—at a cost set by the pole solve, not by the number of variables. The measured cost on the production discretization is a 398.0s pole solve followed by a 370.0s contraction for the full 45-coordinate gradient (Supplementary Section S14).

Boundary form. For an interface Γ between isotropic media moved with normal velocity h , the numerator of (13) reduces to a Hadamard boundary integral,

$$\mathbf{v}^T (\partial_p \mathbf{A}) \mathbf{u} \propto \int_{\Gamma} h \left[(\epsilon_1 - \epsilon_2) \mathbf{E}_{L\parallel} \cdot \mathbf{E}_{R\parallel} - (\epsilon_1^{-1} - \epsilon_2^{-1}) D_{L\perp} D_{R\perp} \right] dA, \quad (14)$$

in which tangential $\mathbf{E}$ and normal $\mathbf{D}$ appear because those are the components continuous across a sharp dielectric boundary. Equation (14) is what makes hole radii, hole positions, and sidewalls differentiable in one framework: all are pieces of the same movable boundary. The proportionality constant and sign depend on the exact definition of $\mathbf{A}$ and on the orientation of the normal, and are pinned in this work by tests—translate one planar interface by a known amount, compare with a centred difference on identical connectivity, reverse the normal and confirm the sign flips—rather than by prose.

Mode identity. Equation (13) is a statement about one isolated pole, so every accepted step must certify that the pole it followed is the same mode. Frequency-nearest tracking is insufficient: a small shape step can legitimately move ω_r by many of the old linewidths, and what must remain continuous is the mode, not a frequency window. The optimizer ranks candidates by (i) biorthogonal field overlap with the accepted QNM on a common representation, (ii) localization in the cavity region, (iii) symmetry and polarization sector, (iv) agreement with the first-order pole prediction, and only then (v) frequency proximity as a tie-breaker. A step whose candidates clear none of the overlap and localization gates is declared unresolved and is not allowed to update the quasi-Newton history, because a silent mode switch can make a correct positive directional derivative appear to produce a negative objective change.

Finally, the raw shape gradient is mapped through a boundary metric $(I - \ell^2 \partial_s^2)h = g_Q$ before it is applied, where s is arclength and ℓ a fabrication-aware smoothing length. This is a Sobolev gradient, not an added penalty term. A line search is only admissible if the applied geometry satisfies $\text{dist}(\Gamma(\alpha), \Gamma(0)) \rightarrow 0$ as $\alpha \rightarrow 0$; an early revision of the multilayer pipeline violated this through polygon repair and remeshing, and is retained as a documented counterexample in Supplementary Section S18.

4 What is validated, and how

Analytical self-consistency and physical transfer are separate gates, and the distinction is load-bearing throughout this paper. Analytical Q_{ana} and FDTD Q_{FDTD} are treated as two different observables; no statement below converts one into the other.

4.1 Operator-level gates

At the analytical level we require, at every accepted geometry: a passive pole ($\text{Im } \tilde{\omega} < 0$); an operator residual $\|\mathbf{A}\mathbf{u}\|$ at the solve tolerance; complex symmetry of $\mathbf{A}$ to the level predicted by reciprocity; convergence of $\tilde{\omega}$ under independent refinement of the Zernike order, the vertical order, and the contour quadrature; agreement of the analytic gradient (13) with centred finite differences on identical connectivity; and clearance of the mode-identity gates of Section 3. Representative values for the production discretizations are an operator residual of 2.29×10^{-10} , reciprocity residuals between 1.1×10^{-16} and 3.8×10^{-16} , left/right null-vector equation errors below 2.35×10^{-15} , and agreement between the analytic vector-Jacobian product and a central difference of the Q derivative to 6.71×10^{-6} relative (Supplementary Section S14). These checks certify that the implementation differentiates its own truncation. They say nothing about whether that truncation's absolute Q equals a converged Maxwell solution.

Table 1: Ringdown adequacy of the four quoted endpoints, computed from the retained simulation metadata. $\mathcal{N} < 1$ means the fit observed less than one amplitude lifetime. Design23 diagnostics do not retain a source start time, so their windows are upper bounds and their $\mathcal{N}$ is optimistic. Source: `results/ringdown-adequacy.json`.

endpoint	Q_{FDTD}	$1/\gamma$ (ps)	window (ps)	$\mathcal{N}$
Quan baseline	15 449.47	24.64	22.65	0.919
Quan Phase-58	59 337.16	94.5	22.65	0.24
Design23 step 0	18 991.94	15.98	12	0.751
Design23 step 105	58 240.27	47.65	64	1.34

4.2 Physical transfer, and the ringdown condition

Physical claims use frozen endpoint geometries rendered into an independent FDTD solver, with the same source, monitors, domain, mesh, absorber, run time, and resonance-fitting code at both ends of a comparison. FDTD is never called by the optimizer.

A quality factor extracted from a time-domain ringdown is only constrained if the recorded window is long compared with the amplitude lifetime of the mode being fitted. With $\gamma = \omega/(2Q)$ the amplitude decay rate, the relevant figure of merit is the number of lifetimes observed, $\mathcal{N} = (t_{\text{run}} - t_{\text{start}})\gamma$. We audited all 35 retained FDTD observations in this way (Table 1); 24 of them, including three of the four endpoints that carry the quoted gains, were fitted over less than one lifetime.

This single number replaces the vaguer hedging that a project record accumulates. It says exactly what is uncertain and in which direction: a sub-lifetime exponential fit is poorly conditioned and, empirically in this data set, biased high. The clearest internal evidence is that a 12 ps diagnostic partway along the second device’s trajectory reports $Q_{\text{FDTD}} = 156\,520$ at step 61, whereas the 64 ps analyses that follow it report values near 58 240.27 for geometries with a higher analytical Q (Fig. 5). We therefore treat both device gains as provisional and quote them only together with Table 1.

4.3 Predeclared promotion panel

Promotion from provisional to validated requires a matched convergence panel that was declared before any of it was run and is recorded in `results/publication-revalidation-plan.json` (Table 2). Each endpoint is evaluated over the full factorial of four binary axes, giving 16 variants per endpoint and 64 tasks in total. The gain is promoted only if the same mode is identified in every variant, every fit passes its diagnostics, and the optimized Q exceeds the initial Q in *every* executed variant; the central value is then chosen by the predeclared refinement rule. Unreliable fits are retained and marked, never silently dropped. The panel is presently predeclared not submitted, and no build

Table 2: Predeclared matched convergence panel. Full factorial over four binary axes per endpoint, 64 tasks total, with nine required artifacts per task. Source: `results/publication-revalidation-plan.json`.

axis	Quan–Lončar	Design23
steps per wavelength	28 / 35	18 / 20
run time (ps)	60 / 90	64 / 96
PML layers	12 / 16	12 / 16
padding	1.25 / 1.5625	1.0 / 1.25

command in the source repository is able to submit it.

The audit of Table 1 also revises the panel itself. At the Phase-58 amplitude lifetime of 94.5 ps, the predeclared 60 ps and 90 ps run times correspond to only $\mathcal{N} = 0.63$ and 0.95 lifetimes. The declared axis is thus not by itself sufficient to certify a Q of that size, and the promotion run must extend the run-time axis until $\mathcal{N} \gtrsim 3$ at the optimized endpoint. We record this as a required amendment rather than quietly reinterpreting the original declaration.

5 Device I: air-clad Quan–Lončar nanobeam

The first device is the canonical deterministic-design nanobeam cavity [11]: a mirror-symmetric silicon beam, 700 nm × 220 nm in cross-section, refractive index 3.46, suspended in air, with 15 elliptical through-holes per side and a resonance near 1 502.32 nm. The free variables are the 15 positive-side gaps and the two semi-axes of each hole, i.e. 45 coordinates. Because the cladding is homogeneous, the reference Green tensor is the one of Eq. (7) with a single air cladding, and the production discretization is a $P2$ hole basis with one vertical order against a $y3$ – $z1$ waveguide cross-section basis.

What the optimizer did. The accepted design is the third step of a constrained ascent in which each step was a scaled L^2 displacement of at most 1 nm, with a maximum single-coordinate move of about 0.34 nm. Reconstructing both endpoint geometries and reducing them to per-hole descriptors (Fig. 3, upper row) shows that this is *local refinement* and nothing more: the mean in-plane aspect ratio moves from 1 to 1.0086, with a maximum of 1.016 near the seventh hole; the largest transverse semi-axis change over all holes is 0.819 nm; the transverse air fraction $2r_y/W$ moves from a maximum of 0.3465 to 0.3482; and the mean hole area changes by a factor of 0.9937. The only large motion is longitudinal: hole centres move by up to 24.6 nm, outward at the mirror end and inward near the centre, i.e. a mild stretch of the taper envelope. The device family is unchanged.




Figure 2: Matched-axis top views of the frozen Quan-Lončar endpoints. Silicon is blue, apertures white. Both panels are reconstructed from the retained seed files (`quan_baseline.json`, `quan_optimized_phase58.json`); no field array enters the drawing. The two geometries are visually almost identical, which is the result: see Fig. 3 for the per-hole differences.

Result. Under matched 24 ps independent FDTD analyses of the two frozen geometries, the fitted quality factor rises from 15 449.47 at 1 502.32 nm to 59 337.16 at 1 499.97 nm, a factor of 3.8407. The corresponding analytical values are 7 455.16 and 18 037.3, a factor of 2.42: the surrogate under-predicts the absolute FDTD quality factor at both ends, by factors of 0.483 and 0.304 respectively, and under-predicts the gain as well.

Condition on this number. Both FDTD analyses used identical sources, monitors, domain, mesh, absorber and run time, so the *ratio* is a matched comparison. The absolute values are not converged: at $Q_{\text{FDTD}} = 59\,337.16$ the amplitude lifetime is 94.5 ps while the usable ringdown window was only 22.65 ps, i.e. $\mathcal{N} = 0.24$ lifetimes (Table 1). A fit over a quarter of a lifetime constrains the decay rate weakly and, in this data set, biases it low—hence Q high. This gain is therefore reported as a provisional candidate pending the panel of Table 2 with the run-time amendment noted there.

A separate eight-step continuation of the same ascent reached an analytical Q of 30 307.26 but was never evaluated in FDTD. It is retained in the record under a distinct name and is not part of any claim here.

6 Device II: multilayer Design23 cavity

The second device tests whether the construction survives a substantially harder reference problem. The stack is, from the bottom: an SiO_2 substrate half-space ($n = 1.45$); a 300 nm layer containing the patterned Si_3N_4 strip ($n =$

2.0) with air elsewhere; 200 nm anthracene ($n = 1.8$); 200 nm PVA ($n = 1.5$); and an air half space. The strip is 556.264 nm wide and carries 42 elliptical through-holes per side, giving 126 free coordinates. The resonance sits near 792.347 nm.

Three ingredients are needed here that the first device does not require. The reference symbol is stratified rather than homogeneous, built from the Fresnel recursion and vertical kernels of Section 2; because the stack breaks z reflection, the symmetry projector retains a fixed (x, y) sector while carrying both z parities. The finite-width strip is then added by a compact Dyson step, and the E_y -dominant guided feed mode appears as a simple pole β_f of that step. Its Laurent terms,

$$\mathbf{G}_{\text{wg}}(\beta) = \frac{R_+}{\beta - \beta_f} + \frac{R_-}{\beta + \beta_f} + \text{regular}, \quad R_- = -R_+^T, \quad (15)$$

are subtracted before the longitudinal quadrature and restored exactly in real space as ordered right- and left-going feedthrough propagation, with a factor one half at coincident nodes. Attempting the contour integral without this subtraction is not merely inaccurate: the finite-hole self term alone changes the guided block by about 8 % (Supplementary Section S10).

Trajectory. A pure BFGS ascent on $\log Q$ accepted 105 of 106 attempted steps, raising the analytical quality factor from 8 958.74 to 717 003, a factor of 80.03. Independent FDTD diagnostics were taken every three steps (Fig. 5); they do not follow that scale. The exact retained endpoints rise from 18 991.94 at step 0 to 58 240.27 at step 105, a factor of 3.0666. The initial point is the step-0 endpoint of *this* trajectory; a Q of 20 942.06 belonging to a




Figure 3: Per-hole shape descriptors before and after optimization, plotted from the cavity centre outward. Note the vertical scales: the Quan-Lončar aspect ratio moves by about one percent (upper left) whereas the Design23 aspect ratio moves by up to a factor of 3.112 (lower left), and the Design23 transverse air fraction reaches 0.5938 of the beam width (lower right). The two campaigns therefore exercise different optimization regimes: local refinement and device restructuring. Source: `results/geometry-metrics.json`.

separately qualified seed lineage exists in the record and is deliberately not substituted for it.

One diagnostic, at step 70, reports $Q_{\text{FDTD}} = 436.3$ with a resonance-fit error of 0.0725, some two orders of magnitude above the surrounding scatter, together with a wavelength excursion of nearly 30 nm. It is excluded from every aggregate statistic by a fit-quality rule declared in advance—resonance error above 0.01—and is retained and marked in Fig. 5 rather than deleted. It is the only exclusion; the remaining 32 of 33 observations are used as recorded.

What the optimizer did. Unlike the first device, this is a restructuring of the cavity rather than a refinement of it (Figs. 3 and 4, lower row). The mean longitudinal semi-axis contracts from 76.157 nm to 45.96 nm while the mean transverse semi-axis expands from 76.166 nm to 118.16 nm, reaching 165.1 nm. The mean in-plane aspect ratio therefore moves from 1.0001—circular to four decimal places—to 2.5613, with a maximum of 3.112. Near-circular holes become transverse slots spanning up to 0.5938 of the beam width, leaving dielectric rails of 113 nm on each side, down

from 187.6 nm. The mean hole *area* changes by only a factor of 0.9501: the optimizer reshaped the air inclusions at nearly constant air filling rather than removing or adding material. The mirror section also contracts, the outermost hole centre moving inward by 268 nm through the accumulation of small pitch reductions.

Spectral and modal signature. The restructuring is accompanied by a blueshift of 21.8 nm (792.347 nm $\rightarrow$ 770.497 nm) and by a halving of the effective mode volume, from 5.047 to 2.652 $(\lambda/n)^3$, a factor of 0.525. A concentration of field energy into the air slots at fixed period would produce exactly this pair of signatures, and the geometry change is what one would design if one wanted an air-type rather than a dielectric-type mode. We report this as a *characterized geometric and spectral restructuring consistent with* a change of mode character, and explicitly not as a proven band-edge swap: the present evidence is a correlation of hole shape, wavelength and mode volume, and the diagnostic that would settle it—a band structure of the terminal mirror cell, or a decomposition of the retained




Figure 4: Matched-axis Design23 geometry before and after 105 accepted analytical steps. Si_3N_4 is blue, apertures white. The source is the pair of retained 128-entry geometry vectors (126 free coordinates plus the two fixed cross-section dimensions).

field energy between the dielectric and the air regions—was not part of this campaign. A useful counterpoint is that $\lambda/2p$ evaluated on the mean pitch is 1.689 at step 0 and 1.686 at step 105: the ratio of wavelength to period is essentially unchanged, so arithmetic on the period alone cannot establish a band swap either way.

Condition on this number. The two endpoints were *not* analyzed under matched settings: step 0 used a 12 ps ringdown and step 105 used 64 ps, i.e. $\mathcal{N} = 0.751$ and 1.34 lifetimes respectively. The mismatch is in the direction that would *inflate* the earlier value and therefore *deflate* the quoted gain, but it is a mismatch and the gain is provisional for that reason in addition to the panel requirement.

7 Directions transfer; magnitudes do not

The two campaigns provide an unusually direct test of what a truncated open-system model is good for, because the surrogate errs in *opposite directions* on the two devices and, more strikingly, in opposite directions at the two ends of a single trajectory.

Magnitudes do not transfer. On the first device the surrogate under-predicts the FDTD quality factor at both endpoints, by factors of 0.483 and 0.304. On the second it begins by under-predicting, $Q_{\text{ana}}/Q_{\text{FDTD}} = 0.472$ at step 0, and ends by over-predicting by more than an order of magnitude, 12.31 at step 105. Across the 32 qualified observations the ratio spans 0.472 to 12.98, a range of $27.5\times$ (Fig. 6, right). No single calibration constant, and no monotone reparameterization of one observable into the

Table 3: Agreement between the analytical surrogate and independent FDTD on the Design23 trajectory. Spearman ρ tests monotone ordering without assuming a shared scale; the sign test asks whether consecutive diagnostics move the same way, with a two-sided binomial p against a fair coin. Source: `results/surrogate-fidelity.json`.

	all qualified	12 ps	64 ps
points n	32	23	9
Spearman ρ	0.774	0.911	0.417
sign agreement	19/31	14/22	5/8
sign test p	0.28	0.29	0.73
$Q_{\text{ana}}/Q_{\text{FDTD}}$ span	$27.5\times$	$6.36\times$	$1.17\times$

other, is available. Any procedure that used the analytical Q as an estimate of the physical Q —for reporting, for stopping, or for accepting a step—would have been badly misled, in both directions, by this data.

Part of that drift is a change in the FDTD observable rather than in the surrogate: the protocol changed from 12 ps to 64 ps between steps 67 and 80, and the Q values that a short ringdown assigns to a long-lived mode are biased high (Section 4). Splitting the trajectory at that boundary is therefore mandatory, and we report every aggregate for the two segments separately as well as together (Table 3). The bias drift survives the split: it runs from 0.472 to 3 within the 12 ps segment alone, before the protocol changes at all.

Ordering does transfer, at trajectory scale. Within the 12 ps segment—23 observations under one fixed FDTD protocol—the rank correlation between the two observables is $\rho = 0.911$. The surrogate therefore orders designs along the trajectory in close agreement with the independent solver even while mis-stating every indi-




Figure 5: The 105-step trajectory. Top: analytical and independent FDTD quality factor on a logarithmic axis. Middle and bottom: effective mode volume and resonance wavelength. The dashed line marks the change of FDTD protocol from 12 ps to 64 ps; the open red marker is the single observation excluded by the predeclared fit-quality rule. Analytical and FDTD Q are distinct observables plotted on a shared axis for comparison of *trend*, not of value.

vidual value by a drifting factor. This, and not any absolute agreement, is what makes the model useful as an optimization objective.

But not step by step. We state the limit of this claim as precisely as the data allows. Consecutive diagnostics are three optimizer steps apart, and asking whether the two observables move in the same direction between them gives 19 agreements out of 31 comparisons—61 %, with a two-sided binomial p of 0.28 against a fair coin. Restricting to the larger moves does not help. The per-step sign agreement in this data set is therefore *not* distinguishable from chance, and we do not claim it. The direction-transfer evidence is the rank correlation over the trajectory together with the two matched endpoint comparisons, not a step-level correspondence: small analytical improvements are not resolvable against FDTD fit scatter, and a user of this method should expect to verify progress over tens of steps rather than one.

A sharper historical instance. The clearest single demonstration that the correct treatment of the continuum is what carries direction transfer comes from a repair to the second device’s solver. A saved geometry direction was evaluated three ways: independent FDTD reported an endpoint ratio of 1.0007628—flat; the pre-repair an-




Figure 6: Left: analytical against independent FDTD quality factor on matched logarithmic axes. The diagonal is a reference line, not a model—the two axes carry different observables. Right: the ratio $Q_{\text{ana}}/Q_{\text{FDTD}}$ against optimizer step, showing that the surrogate bias is not a constant calibration factor but drifts by 27.5 $\times$ along a single trajectory. Source: `results/surrogate-fidelity.json`.

alytical operator reported 3.6045899, an apparent 3.6 $\times$ improvement; and the repaired pole-subtracted operator reported 1.0000101, correctly flat. The surrogate’s value of a direction can be wrong by a factor of three and a half when the feed pole is mishandled, and it is the operator construction of Section 2—not a fitted correction—that makes it right.

8 Discussion

The practical conclusion of this work is narrower than numerical agreement between solvers, and more useful. A reduced outgoing model can supply productive ascent directions in a high-dimensional geometry space while its absolute linewidth is wrong by an order of magnitude—and while the sign and size of that error drift along the very trajectory it is driving. The two roles should therefore be separated by construction: analytical convergence certifies the local reduced problem and supplies the gradient, and a sparse independent solver certifies transfer and determines every reported physical number. Our results support that division and would not support any weaker one.

This has a concrete consequence for how such an optimization is run. Because the surrogate bias drifts, a stopping criterion or a step-acceptance test phrased in absolute analytical Q is unsafe; because per-step sign agreement with an independent solver is not resolvable (Section 7), a periodic independent check placed every few steps cannot certify individual steps either. What the evidence does support is independent verification at trajectory scale—at the endpoints, under matched settings, with a ringdown long compared with the fitted lifetime.

The two devices probe genuinely different failure modes, which is why running both matters more than optimizing either one further. In the air-clad cavity, basis order and the weak radiation channels dominate the error budget. In the multilayer cavity, the stratified continuum quadrature, the feed-pole subtraction, mode tracking through a

large geometry change, and geometry transfer to an independent mesh are all additional and independent sources of error; the historical instance in Section 7, in which a mishandled feed pole turned a flat direction into an apparent $3.6\times$ gain, shows how consequential that machinery is. Success on both is evidence for operator-level portability rather than for a second optimization of the same device.

The two campaigns also exercise different optimization regimes, and the contrast is informative in itself. The first device was refined: aspect ratios moved by one percent and the device family was preserved. The second was restructured: circular holes became transverse slots at nearly constant air area, the mode volume halved, and the resonance blueshifted by 21.8 nm. That a gradient built from an analytically continued continuum can carry a design that far from its seed—while an independent solver still confirms a gain—is the more surprising of the two outcomes, and it is also the one whose mode-character interpretation we have deliberately left open pending a diagnostic that can settle it.

What this work does not establish should be equally clear. It does not establish fabrication robustness, experimental quality factors, or agreement of the gradients themselves between the analytical and FDTD descriptions; each of those is a separate study. It does not establish converged absolute quality factors for either device: the panel of Table 2, amended for run time as described, is required for that, and until it is executed the two gains remain **PROVISIONAL** in the accompanying claims ledger. And it does not establish that the second device’s mode changed character, only that its geometry, wavelength, and mode volume all moved together in the way such a change would produce.

9 Methods summary

Analytical poles are located by continuation from a known symmetry sector along the frozen contour homotopy of Eq. (8) and accepted only when the operator residual, the passive sign convention, complex symmetry, and the branch-overlap test all pass. Gradients are validated against both coordinate and random-direction central differences at the production discretization; the representative agreement is 6.71×10^{-6} relative on the Q derivative. Geometry optimization uses bound-constrained quasi-Newton updates on $\log Q$ with a trust region, a Sobolev boundary metric, and a line search that rejects any candidate whose branch overlap falls below 0.999. Exact thresholds, modal orders, quadrature rules, and optimizer settings are tabulated in Supplementary Sections S13 and S14.

FDTD is never called by the optimizer and never supplies a gradient. Frozen endpoint geometries are rendered into independent simulations and analyzed with the same resonance-fitting code at both ends of any comparison. The publication protocol, including the predeclared factorial panel and its promotion gate, is enumerated in Sup-

plementary Section S15. No historical fit is promoted because it agrees with the analytical trend, and no analytical value is substituted for a physical one.

Data and code availability

The manuscript source, evidence tables, figure generators, complete derivation archive, and implementation snapshot are maintained in a provenance-preserving repository. Every number printed in this paper and its Supplementary Information is generated from a machine-readable record under **results/** and injected into the $\text{\LaTeX}$ source by a single exporter, so that a replaced measurement changes the document without any prose being retyped. Raw field arrays are retained outside version control and indexed by absolute path, size, and SHA-256 checksum. Reviewer and archival access, together with the released subset, will be specified at submission.

Acknowledgments

Author list, affiliations, funding, and computational-resource statements are in preparation. The authors will disclose the role of automated coding and language tools in accordance with the target venue’s policy.

References

- [1] Jakob S. Jensen and Ole Sigmund. Topology optimization for nano-photonics. *Laser & Photonics Reviews*, 5(2):308–321, 2011.
- [2] Sean Molesky, Zin Lin, Alexander Y. Piggott, Weiliang Jin, Jelena Vuckovic, and Alejandro W. Rodriguez. Inverse design in nanophotonics. *Nature Photonics*, 12:659–670, 2018.
- [3] Tyler W. Hughes, Momchil Minkov, Ian A. D. Williamson, and Shanhui Fan. Adjoint method and inverse design for nonlinear nanophotonic devices. *ACS Photonics*, 5(12):4781–4787, 2018.
- [4] Momchil Minkov, Ian A. D. Williamson, Lucio C. Andreani, Dario Gerace, Beicheng Lou, Alex Y. Song, Tyler W. Hughes, and Shanhui Fan. Inverse design of photonic crystals through automatic differentiation. *ACS Photonics*, 7(7):1729–1741, 2020.
- [5] Momchil Minkov, Vincenzo Savona, and Dario Gerace. Photonic crystal slab cavity simultaneously optimized for ultra-high Q/V and vertical radiation coupling. *Applied Physics Letters*, 111(13):131104, 2017.
- [6] Lucio Claudio Andreani and Dario Gerace. Photonic-crystal slabs with a triangular lattice of triangular holes investigated using a guided-mode expansion method. *Physical Review B*, 73:235114, 2006.

- [7] Philippe Lalanne, Wei Yan, Kevin Vynck, Christophe Sauvan, and Jean-Paul Hugonin. Light interaction with photonic and plasmonic resonances. *Laser & Photonics Reviews*, 12(5):1700113, 2018.
- [8] Christophe Sauvan, Jean-Paul Hugonin, Ivan S. Maksymov, and Philippe Lalanne. Theory of the spontaneous optical emission of nanosize photonic and plasmon resonators. *Physical Review Letters*, 110:237401, 2013.
- [9] Philip Trost Kristensen, Kathrin Herrmann, Francesco Intravaia, and Kurt Busch. Modeling electromagnetic resonators using quasinormal modes. *Advances in Optics and Photonics*, 12(3):612–708, 2020.
- [10] Philippe Lalanne, Wei Yan, Alexandre Gras, Christophe Sauvan, Jean-Paul Hugonin, Mondher Besbes, Guillaume Demesy, M. D. Truong, Boris Gralak, Frederic Zolla, Andre Nicolet, Felix Binkowski, Lin Zschiedrich, Sven Burger, Jorn Zimmerling, Rob Remis, Paul Urbach, H. T. Liu, and Thomas Weiss. Quasinormal mode solvers for resonators with dispersive materials. *Journal of the Optical Society of America A*, 36(4):686–704, 2019.
- [11] Qimin Quan and Marko Lončar. Deterministic design of wavelength scale, ultra-high Q photonic crystal nanobeam cavities. *Optics Express*, 19(19):18529–18542, 2011.

Supplementary Information

Inverse design of open nanophotonic cavities through differentiable mode expansions

Author list and affiliations in preparation

September 5, 2026

This document is self-contained. Sections S1–S13 develop the theory and the numerical construction in one consistent notation, starting from Maxwell’s equations and ending at the optimizer. Sections S14–S19 present the evidence: analytical convergence and gradient validation, the independent-FDTD protocol and its ringdown audit, the two device campaigns in full, a catalogue of rejected approximations and outright failures, and a reproducibility and provenance appendix that maps every theory section onto the retained derivation record. Every numerical value printed here is generated from a machine-readable record; none is transcribed by hand.

Contents

S1	Scope, notation, and the claim ledger	4
	S1.1 What this document contains	4
	S1.2 Two observables, never interchanged	4
	S1.3 Claim ledger	4
	S1.4 Notation	5
S2	Conventions, quasinormal modes, and normalization	5
	S2.1 Time convention and the pole	5
	S2.2 Nonlinear eigenproblem and biorthogonality	5
	S2.3 Normalization	6
	S2.4 The logarithmic objective	6
S3	Exact parameterized dielectric and its transforms	6
	S3.1 Reference and contrast	6
	S3.2 Closed-form Fourier factors	7
	S3.3 Analytic geometry derivatives	7
	S3.4 What this parameterization buys	8
S4	Outgoing integral equation and the regularized determinant	8
	S4.1 Reduction to the holes	8
	S4.2 Why the determinant must be regularized	8
	S4.3 Finite sequence	9
	S4.4 Relation to cell scattering matrices	9
S5	Reference Green tensor: homogeneous cladding	9
	S5.1 Fourier reduction along the beam	9
	S5.2 Exact homogeneous symbol	10
	S5.3 Cross-section basis and Galerkin matrices	10
	S5.4 Dyson resummation of the core	10
	S5.5 Guided poles are not inserted by hand	11
	S5.6 On-shell evaluation inside the light cone	11
S6	Reference Green tensor: stratified stack	11

	S6.1	Why the homogeneous reference fails here	11
	S6.2	Generalized Fresnel recursion	12
	S6.3	Scalar vertical kernels and the Wronskian	12
	S6.4	Adding the strip	12
S7		The outgoing contour and its quadrature	13
	S7.1	Branch topology in scaled momentum	13
	S7.2	The deformation	13
	S7.3	Orientation test	13
	S7.4	Panelized quadrature	13
	S7.5	Known failure modes	14
S8		Aperture basis and projection	14
	S8.1	Normalized hole coordinates	14
	S8.2	Zernike–Legendre vector basis	14
	S8.3	Exact thickness projection	15
	S8.4	Projection into the cross-section basis	15
	S8.5	Transpose, not conjugate	15
	S8.6	Production orders	15
S9		Symmetry reduction and finite-array assembly	16
	S9.1	The assembled operator	16
	S9.2	Symmetry sectors	16
	S9.3	Finite array, not a periodic cell	16
	S9.4	Block structure and cell cascades	16
S10		Feed-pole extraction and residues	16
	S10.1	The problem	16
	S10.2	Laurent extraction	17
	S10.3	Exact real-space restoration	17
	S10.4	Magnitude of the effect	17
	S10.5	What must be differentiated	18
S11		Continuation, pole search, and mode identity	18
	S11.1	Locating the first pole	18
	S11.2	Continuation off the real axis	18
	S11.3	Mode identity under geometry change	18
	S11.4	Multiple poles and near-degeneracies	19
S12		Pole and quality-factor derivatives	19
	S12.1	The implicit derivative	19
	S12.2	Assembling $\partial_{\theta_j} \mathbf{A}$	19
	S12.3	Reverse-mode contraction	20
	S12.4	Boundary form	20
	S12.5	Boundary metric	20
	S12.6	Trust ladder	21
S13		Optimizer, trust region, and acceptance gates	21
	S13.1	Objective and variables	21
	S13.2	Step control	21
	S13.3	Acceptance gates	22
	S13.4	One loop, written out	22
	S13.5	What the independent solver is not allowed to do	22
S14		Analytical convergence and gradient validation	23
	S14.1	Operator identities	23
	S14.2	Refinement convergence	23
	S14.3	Gradient validation	23

	S14.4	A directional gate	24
	S14.5	Complete audited value table	24
S15		Independent-FDTD protocol and the ringdown audit	25
	S15.1	Role and separation	25
	S15.2	Endpoint procedure	25
	S15.3	Ringdown adequacy	26
	S15.4	Predeclared promotion panel	26
	S15.5	A required amendment to the panel	27
	S15.6	Complete adequacy table	27
S16		Air-clad Quan–Lončar campaign	27
	S16.1	Device and parameterization	27
	S16.2	A lineage caution	28
	S16.3	Optimization	28
	S16.4	Geometry outcome	28
	S16.5	Independent evaluation	28
	S16.6	Status	29
S17		Multilayer Design23 campaign	29
	S17.1	Device and parameterization	29
	S17.2	Trajectory	29
	S17.3	The excluded observation	30
	S17.4	Protocol change	30
	S17.5	Geometry outcome	31
	S17.6	Spectral and modal signature	31
	S17.7	Independent evaluation and status	31
S18		Rejected approximations and failures	32
	S18.1	Few-mode and closed bases	32
	S18.2	Unregularized collocation determinant	32
	S18.3	Real-axis quadrature at complex frequency	32
	S18.4	Hermitian conjugate in place of transpose	32
	S18.5	Layered cavity without feed-pole subtraction	32
	S18.6	The flat direction that looked like a $3.6\times$ gain	33
	S18.7	Mesh-motion blockade	33
	S18.8	Short-ringdown quality factors	33
	S18.9	Objectives that have not passed their own gates	33
S19		Reproducibility and provenance	34
	S19.1	How the numbers reach this document	34
	S19.2	Statistical rules declared in advance	34
	S19.3	Provenance of the theory sections	34
	S19.4	Raw artifacts	35
	S19.5	Build	35

Theory and construction

S1 Scope, notation, and the claim ledger

S1.1 What this document contains

The main text states a construction and two device results in the space available to it. This document gives the construction in full, in one notation, and gives the evidence in full, including the parts that failed.

Sections S2–S13 are a linear development: conventions and the quasinormal-mode (QNM) normalization (§S2); the exact parameterized dielectric and its closed-form transforms (§S3); the outgoing volume-integral equation and its regularized determinant (§S4); the reference Green tensor for a homogeneous cladding (§S5) and for a stratified stack (§S6); the outgoing contour and its quadrature (§S7); the aperture basis (§S8); symmetry reduction and array assembly (§S9); feed-pole extraction (§S10); complex-frequency continuation and mode identity (§S11); the pole and Q derivatives (§S12); and the optimizer (§S13).

Sections S14–S19 are evidence: analytical convergence and gradient validation (§S14); the independent-FDTD protocol and the ringdown audit (§S15); the two device campaigns (§S16, §S17); a catalogue of rejected approximations and failures (§S18); and reproducibility with a provenance map (§S19).

S1.2 Two observables, never interchanged

Throughout, Q_{ana} denotes a quality factor computed from a pole of the analytical operator of this work, and Q_{FDTD} denotes a quality factor fitted to the ringdown of an independent finite-difference time-domain simulation. They are different observables computed by different means from different approximations of the same device. No expression in this work converts one into the other, and no plot places them on a shared axis without saying so. Where a ratio $Q_{\text{ana}}/Q_{\text{FDTD}}$ appears it is a measure of disagreement between two models, not an error in either.

S1.3 Claim ledger

Every quantitative statement in the main text carries one of three statuses, recorded in a machine-readable ledger alongside the source.

validated

The construction of the open-pole operator (§S4–§S9) and its differentiability (§S12). The supporting evidence is §S14: operator residual, complex symmetry, refinement convergence, and analytic-versus-finite-difference gradient agreement.

provisional

The two device gains, $3.8407\times$ and $3.0666\times$, and the claim that both devices exhibit transfer. These are historical endpoint measurements under stated FDTD settings. They become validated only when the predeclared panel of §S15 passes with the run-time amendment recorded there.

retired

Results superseded during development and retained only as counterexamples, catalogued in §S18. The most consequential is a pre-repair analytical operator that reported a $3.6045899\times$ improvement along a direction that independent FDTD showed to be flat.

S1.4 Notation

symbol	meaning
$\tilde{\omega} = \omega_r + i\omega_i$	complex pole, $\omega_i < 0$
$k_0 = \omega/c$, $k_c = n_c\omega/c$	vacuum and cladding wavenumbers
β	longitudinal (beam-axis) momentum
$u = \beta/k_c$	scaled longitudinal momentum
ϵ_b, ϵ_c	core and cladding permittivity
$\chi_b = \epsilon_b - \epsilon_c$	core susceptibility contrast
W, T	beam width and patterned-layer thickness
$x_j, r_{x,j}, r_{y,j}$	centre and semi-axes of hole j
$\boldsymbol{\theta}$	vector of free geometry coordinates
$\mathbf{G}_c, \mathbf{G}_{wg}, \mathbf{G}_b$	cladding, waveguide, stratified Green tensors
$\mathbf{G}_H^+$	outgoing Green matrix projected on the holes
$\mathbf{A}(\omega, \boldsymbol{\theta})$	assembled pole operator
$\mathbf{u}, \mathbf{v}$	right and left null vectors of $\mathbf{A}$
N_{xy}, N_z, N_β	aperture, vertical, and contour orders
$\mathcal{N}$	FDTD ringdown windows per amplitude lifetime

Vectors are bold, matrices upright bold, and the transpose $(\cdot)^T$ is used wherever analytic continuation is required; the Hermitian conjugate appears only where a quantity is evaluated on the real-frequency boundary and is explicitly flagged there.

S2 Conventions, quasinormal modes, and normalization

S2.1 Time convention and the pole

All fields carry $\exp(-i\omega t)$. Maxwell's equations for nonmagnetic, nondispersive, isotropic media reduce to

$$\nabla \times \nabla \times \mathbf{E} - \frac{\omega^2}{c^2} \epsilon(\mathbf{r}) \mathbf{E} = 0. \quad (\text{S1})$$

A resonance of an open system is not an eigenvalue of (S1) under square-integrable boundary conditions; it is a pole of the associated scattering problem, obtained by imposing purely outgoing behaviour and continuing ω off the real axis. With the above time convention a passive outgoing pole satisfies

$$\tilde{\omega} = \omega_r + i\omega_i, \quad \omega_i < 0, \quad Q = -\frac{\omega_r}{2\omega_i}. \quad (\text{S2})$$

The sign of ω_i is not a convention to be chosen at the end of a calculation. It is a test: a construction that returns $\omega_i > 0$ for a passive dielectric has either continued onto the wrong Riemann sheet or has an incoming rather than outgoing Green function, and both errors occurred during development (§S18).

The associated field grows exponentially at large distance—the well-known QNM divergence—which is why every normalization used here is defined through the operator rather than through a field-energy integral over an unbounded region [1, 2, 3].

S2.2 Nonlinear eigenproblem and biorthogonality

Every discretization used in this work produces a finite matrix that depends nonlinearly on frequency,

$$\mathbf{A}(\tilde{\omega}, \boldsymbol{\theta}) \mathbf{u} = 0, \quad \mathbf{v}^T \mathbf{A}(\tilde{\omega}, \boldsymbol{\theta}) = 0. \quad (\text{S3})$$

The nonlinearity is essential and not an artifact: the reference Green tensor, the radiation branch cuts, and the contour of §S7 all depend on ω . Equation (S3) is a nonlinear eigenvalue problem [4, 5], and $\mathbf{v}, \mathbf{u}$ are its left and right null vectors at a simple root of $\det \mathbf{A}$.

The pairing between $\mathbf{v}$ and $\mathbf{u}$ is the transpose bilinear form, not the Hermitian sesquilinear form. For real-valued basis functions and real ω the two coincide, which is why the error is easy to make and hard to notice; but only the bilinear form is analytic in ω , so only it survives continuation into the lower half plane. Reciprocity of the underlying Maxwell problem implies that $\mathbf{A}$ is complex symmetric, $\mathbf{A} = \mathbf{A}^T$, whence $\mathbf{v} = \mathbf{u}$ up to normalization. We do *not* impose this. We compute $\mathbf{v}$ and $\mathbf{u}$ independently and use $\|\mathbf{A} - \mathbf{A}^T\|/\|\mathbf{A}\|$ as a running check on the entire assembly chain; measured values are between 1.1×10^{-16} and 3.8×10^{-16} (§S14).

S2.3 Normalization

Let $\tilde{\omega}$ be a simple root. Differentiating (S3) with respect to any parameter p and contracting with $\mathbf{v}$ gives

$$\mathbf{v}^T(\partial_p \mathbf{A})\mathbf{u} + \frac{d\tilde{\omega}}{dp} \mathbf{v}^T(\partial_\omega \mathbf{A})\mathbf{u} = 0, \quad (\text{S4})$$

where the field derivative $\partial_p \mathbf{u}$ has dropped out because $\mathbf{v}^T \mathbf{A} = 0$. The quantity

$$\mathcal{M} = \mathbf{v}^T(\partial_\omega \mathbf{A})\mathbf{u} \quad (\text{S5})$$

is the QNM normalization *for this operator*. It is nonzero exactly when the root is simple, it carries the frequency dependence of the reference Green tensor and the contour, and it is invariant under independent rescaling of $\mathbf{v}$ and $\mathbf{u}$ in the combination that appears in (S4). Everything in §S12 follows from (S4) alone.

Substituting a closed-cavity energy normalization $\int \epsilon |\mathbf{E}|^2$ for $\mathcal{M}$ is the single most common way to get a plausible-looking $\partial\omega_r/\partial p$ together with a wrong $\partial\omega_i/\partial p$. Since Q depends on ω_i and $|\omega_i|/\omega_r = 1/(2Q)$, that error is invisible in the wavelength and fatal in the objective.

S2.4 The logarithmic objective

Writing $\tilde{\omega} = a + ib$ with $b < 0$, (S2) gives

$$\delta Q = -\frac{\delta a}{2b} + \frac{a \delta b}{2b^2}, \quad \delta \log Q = \frac{\delta a}{a} - \frac{\delta b}{b}. \quad (\text{S6})$$

We optimize $\log Q$: multiplicative gains become additive, the two components of the pole enter symmetrically in relative terms, and the conditioning of the problem is exposed rather than hidden. Since $|b|/a = 1/(2Q)$, at $Q = 10^4$ the loss rate is 5×10^{-5} of the resonant frequency. Any part of the construction whose relative accuracy is worse than that—an unresolved continuum quadrature, a contour that has drifted onto another sheet, an absorber reflection—can dominate δb and reverse the sign of the gradient while leaving δa apparently converged. This single inequality motivates most of the numerical care documented in §§S7, S10, and S14.

S3 Exact parameterized dielectric and its transforms

S3.1 Reference and contrast

Let the unpatterned device define the reference permittivity ϵ_0 . For the air-clad beam this is the rectangular core in a homogeneous cladding,

$$\epsilon_0(y, z) = \epsilon_c + (\epsilon_b - \epsilon_c) \mathbf{1}_{|y| < W/2} \mathbf{1}_{|z| < T/2}, \quad (\text{S7})$$

and for the multilayer device it is the planar stack with the unpatterned Si_3N_4 strip embedded in it (§S6). In both cases the patterned device differs from the reference only inside the etched holes.

Hole j is the elliptical through-cylinder

$$H_j(\boldsymbol{\theta}) = \left\{ \mathbf{r} : \frac{(x - x_j)^2}{r_{x,j}^2} + \frac{y^2}{r_{y,j}^2} \leq 1, |z| \leq \frac{T}{2} \right\}, \quad (\text{S8})$$

so that

$$\epsilon_{\theta} = \epsilon_0 + \Delta\epsilon_{\theta}, \quad \Delta\epsilon_{\theta} = (\epsilon_c - \epsilon_b) \sum_{j=1}^{N_h} \mathbf{1}_{H_j}(\theta). \quad (\text{S9})$$

For a mirror-symmetric cavity only the positive- x half is stored,

$$x_1 = \frac{s_0}{2}, \quad x_j = \frac{s_0}{2} + \sum_{\ell=1}^{j-1} s_{\ell}, \quad (\text{S10})$$

with holes at $-x_j$ generated exactly. The free vector is therefore

$$\theta = (s_0, \dots, s_{N_h-1}; r_{x,1}, \dots, r_{x,N_h}; r_{y,1}, \dots, r_{y,N_h}), \quad (\text{S11})$$

of length $3N_h$: 45 coordinates for $N_h = 15$ and 126 for $N_h = 42$. Storing pitches rather than absolute centres makes the mirror-side extent a derived quantity, and is why the optimizer's accumulated pitch reduction on the second device shows up as a 268 nm inward motion of the outermost hole (§S17).

S3.2 Closed-form Fourier factors

With $\hat{f}(\mathbf{k}) = \int f(\mathbf{r}) e^{-i\mathbf{k} \cdot \mathbf{r}} d^3r$, direct integration of (S8) over the elliptical disk and the vertical interval gives, exactly,

$$\hat{\mathbf{1}}_{H_j}(\mathbf{k}) = e^{-ik_x x_j} \underbrace{\left[2\pi r_{x,j} r_{y,j} \frac{J_1(\rho_j)}{\rho_j} \right]}_{\text{elliptical disk}} \underbrace{\left[\frac{2 \sin(k_z T/2)}{k_z} \right]}_{\text{thickness}}, \quad (\text{S12})$$

$$\rho_j^2 = (k_x r_{x,j})^2 + (k_y r_{y,j})^2. \quad (\text{S13})$$

Both bracketed factors have removable singularities, with limits $J_1(\rho)/\rho \rightarrow \frac{1}{2}$ and $2 \sin(k_z T/2)/k_z \rightarrow T$, so that $\hat{\mathbf{1}}_{H_j}(\mathbf{0}) = \pi r_{x,j} r_{y,j} T$ is exactly the hole volume. The implementation evaluates both through small-argument series below a fixed switchover, which is tested against the closed forms to machine precision. Summing (S12) over holes gives the complete geometry dependence of the scatterer,

$$\widehat{\Delta\epsilon_{\theta}}(\mathbf{k}) = (\epsilon_c - \epsilon_b) \sum_j e^{-ik_x x_j} \left[2\pi r_{x,j} r_{y,j} \frac{J_1(\rho_j)}{\rho_j} \right] \left[\frac{2 \sin(k_z T/2)}{k_z} \right]. \quad (\text{S14})$$

This is an exact analytical representation of the entire finite hole array as a function of every centre and every radius. No permittivity is ever sampled onto a grid, and no subpixel smoothing rule enters the model.

S3.3 Analytic geometry derivatives

The derivatives needed by §S12 follow from (S12) in closed form. Translation is a phase,

$$\frac{\partial \hat{\mathbf{1}}_{H_j}}{\partial x_j} = -ik_x \hat{\mathbf{1}}_{H_j}, \quad (\text{S15})$$

and for the radii, writing $f(\rho) = J_1(\rho)/\rho$ with

$$f'(\rho) = \frac{J_0(\rho)}{\rho} - \frac{2J_1(\rho)}{\rho^2}, \quad \frac{\partial \rho_j}{\partial r_{y,j}} = \frac{k_y^2 r_{y,j}}{\rho_j}, \quad (\text{S16})$$

the product rule applied to $2\pi r_{x,j} r_{y,j} f(\rho_j)$ gives

$$\frac{\partial \hat{\mathbf{1}}_{H_j}}{\partial r_{y,j}} = \hat{\mathbf{1}}_{H_j} \left[\frac{1}{r_{y,j}} + \frac{f'(\rho_j)}{f(\rho_j)} \frac{k_y^2 r_{y,j}}{\rho_j} \right], \quad (\text{S17})$$

and symmetrically for $r_{x,j}$ with $k_y \rightarrow k_x$. Because (S10) is affine, the derivative with respect to a pitch s_ℓ is the sum of (S15) over every hole outboard of ℓ ; this is the only place where the geometry parameterization couples coordinates, and it is handled by the reverse-mode accumulation of §S12 rather than by an explicit chain of sums.

For completeness, the rectangular reference contrast of (S7) has the exact transverse transform

$$\widehat{\Delta\epsilon_{\text{beam}}}(k_y, k_z) = (\epsilon_b - \epsilon_c) WT \operatorname{sinc}\left(\frac{k_y W}{2}\right) \operatorname{sinc}\left(\frac{k_z T}{2}\right), \quad (\text{S18})$$

with unnormalized $\operatorname{sinc}(x) = \sin(x)/x$; the beam width W and thickness T are therefore differentiable in the same framework, though both were held fixed in the campaigns reported here.

S3.4 What this parameterization buys

Three properties matter downstream. The model is exact in the geometry variables, so a gradient is a derivative of the true parameterized problem and not of a discretization of it. The derivatives are analytic, so the 6.71×10^{-6} -level agreement with central differences reported in §S14 tests the *implementation* rather than a step-size choice. And the geometry map is continuous as the step size vanishes—no remeshing, no polygon repair, no signed-distance reinitialization—so a line search satisfies $\operatorname{dist}(\Gamma(\alpha), \Gamma(0)) \rightarrow 0$ as $\alpha \rightarrow 0$ by construction. That last property is not free in general and its violation is documented in §S18.

S4 Outgoing integral equation and the regularized determinant

S4.1 Reduction to the holes

Let $\mathcal{L}_0(\omega) = \nabla \times \nabla \times -(\omega^2/c^2)\epsilon_0$ be the reference Maxwell operator and let $\mathbf{G}_{\text{wg}}^+(\omega) = \mathcal{L}_0(\omega)^{-1}$ be its *outgoing* inverse, constructed in §S5–§S7. Splitting (S1) as $\mathcal{L}_0 \mathbf{E} = (\omega^2/c^2)\Delta\epsilon_\theta \mathbf{E}$ and inverting gives the homogeneous Lippmann–Schwinger equation

$$\mathbf{E} = \frac{\omega^2}{c^2} \mathbf{G}_{\text{wg}}^+(\omega) \Delta\epsilon_\theta \mathbf{E}. \quad (\text{S19})$$

Because $\Delta\epsilon_\theta$ vanishes outside $\Omega = \bigcup_j H_j$, the unknown in (S19) may be restricted to Ω : the field everywhere else is recovered by one application of $\mathbf{G}_{\text{wg}}^+$. This is the entire reason the method is cheap. The support of the unknown is a few dozen small ellipses, not a computational domain, and it does not grow when the cavity is made longer or the mirror stronger.

Defining $K(\omega, \theta) = (\omega^2/c^2)\mathbf{G}_{\text{wg}}^+(\omega)\Delta\epsilon_\theta$ restricted to Ω , a nontrivial solution of (S19) exists exactly when $I - K$ has nontrivial kernel, i.e. when

$$D(\omega, \theta) = \det_{\text{reg}}[I - K(\omega, \theta)] = 0. \quad (\text{S20})$$

Together with the passivity requirement $\operatorname{Im} \tilde{\omega} < 0$ and the outgoing choice of sheet, (S20) is an exact implicit definition of $Q(\theta)$:

$$Q(\theta) = -\frac{\operatorname{Re} \tilde{\omega}(\theta)}{2 \operatorname{Im} \tilde{\omega}(\theta)}, \quad D(\tilde{\omega}(\theta), \theta) = 0, \quad \operatorname{Im} \tilde{\omega} < 0. \quad (\text{S21})$$

S4.2 Why the determinant must be regularized

The subscript in (S20) is not decorative. The electric dyadic Green tensor contains a singular local term—in momentum space the longitudinal projector $-P_L/k_c^2$ of (S27), in real space a δ -function plus a nonintegrable $1/R^3$ tail whose principal value depends on the shape of the excluded volume [6]. The operator K is consequently *not* trace class, and the Fredholm determinant $\det(I - K)$ does not exist as a convergent product of eigenvalues.

Three consequences follow.

1. A collocation matrix built by sampling the real-space kernel at quadrature nodes has entries that depend on the exclusion rule chosen at coincident points. Its determinant converges to *something* under mesh refinement, but not to (S20), and calling it exact would be wrong. We do not use one.
2. We keep the longitudinal projector in momentum space, where it is algebraic. The self term is then never evaluated as a singular real-space integral and no depolarization factor is guessed.
3. The determinant is taken in the regularized sense: either as an ordinary Fredholm determinant of the trace-class remainder after the self term is extracted, or as the two-regularized determinant

$$\det_2(I - K) = \det[(I - K)e^K], \quad (\text{S22})$$

which is well defined for Hilbert–Schmidt K [7]. Both have the same zero set as (S20), since e^K is invertible; only the value away from the root differs.

Formally, where the local trace expansions converge,

$$\det(I - K) = \exp\left[-\sum_{m \geq 1} \frac{\text{Tr } K^m}{m}\right], \quad \det_2(I - K) = \exp\left[-\sum_{m \geq 2} \frac{\text{Tr } K^m}{m}\right], \quad (\text{S23})$$

the $m = 1$ term being exactly the divergent self trace that (S22) removes. Equation (S23) identifies the analytic object; it is not how we evaluate it. At a resonance an eigenvalue of K approaches unity and the trace-log series loses its convergence, so the computational realization is always the finite matrix determinant of §S9.

S4.3 Finite sequence

Projecting (S19) onto the truncated interior basis of §S8 gives a finite matrix

$$A_{ab}^{(N)}(\omega, \boldsymbol{\theta}) = \delta_{ab} - \frac{\omega^2}{c^2}(\epsilon_c - \epsilon_b) \int_{\Omega} d^3r \int_{\Omega} d^3r' \phi_a \cdot \mathbf{G}_{\text{wg}}^+ \cdot \phi_b, \quad (\text{S24})$$

whose determinant $D_N(\omega, \boldsymbol{\theta}) = \det A^{(N)}$ has roots $\tilde{\omega}_N$ and finite quality factors $Q_N = -\text{Re } \tilde{\omega}_N / (2 \text{Im } \tilde{\omega}_N)$. The claim is convergence of the joint limit

$$Q(\boldsymbol{\theta}) = \lim_{N_{xy}, N_z, N_{\beta} \rightarrow \infty} Q_{N_{xy}, N_z, N_{\beta}}(\boldsymbol{\theta}), \quad (\text{S25})$$

in which N_{β} is a *quadrature* order for the radiation continuum, not a count of physical radiation modes. Measured convergence under independent refinement of each index is reported in §S14.

S4.4 Relation to cell scattering matrices

Partitioning the indices of (S24) by hole and eliminating interior blocks by Schur complements produces exactly the cell scattering maps that a transfer-matrix or Redheffer-star treatment would cascade. The difference is the direction of the derivation: here the cell matrices are a stable block factorization of a converged global operator, so the radiation channels they omit are known and controllable, whereas a cascade assembled from independently computed cell matrices carries an unknown truncation of the continuum. That distinction is not academic—it is the same failure mode that made an early revision of this work report a numerically precise and physically meaningless Q (§S18).

S5 Reference Green tensor: homogeneous cladding

S5.1 Fourier reduction along the beam

The unpatterned beam is invariant in x , so the reference Green tensor reduces to a one-dimensional inverse transform,

$$\mathbf{G}_{\text{wg}}(x, \mathbf{r}_{\perp}; x', \mathbf{r}'_{\perp}; \omega) = \int_{\mathcal{C}_{\beta}} \frac{d\beta}{2\pi} e^{i\beta(x-x')} \widetilde{\mathbf{G}_{\text{wg}}}(\beta; \mathbf{r}_{\perp}, \mathbf{r}'_{\perp}; \omega). \quad (\text{S26})$$

Everything about openness is carried by the contour $\mathcal{C}_\beta$ (§S7); at fixed β the transverse problem is a compactly supported scattering problem, because only the rectangle $|y| < W/2$, $|z| < T/2$ differs from the cladding.

S5.2 Exact homogeneous symbol

Write $k_c = n_c \omega / c$, $\mathbf{k} = (\beta, q_y, q_z)$ and $k^2 = \mathbf{k} \cdot \mathbf{k}$. For the cladding operator $L_c = \nabla \times \nabla \times -k_c^2 I$, introduce the longitudinal and transverse projectors $P_L = \mathbf{k}\mathbf{k}^T/k^2$ and $P_T = I - P_L$. Then

$$\widetilde{\mathbf{G}}_c(\beta, q_y, q_z; \omega) = \frac{P_T}{k^2 - k_c^2} - \frac{P_L}{k_c^2}, \quad (\text{S27})$$

which is verified directly from $[(k^2 - k_c^2)I - \mathbf{k}\mathbf{k}^T]\widetilde{\mathbf{G}}_c = I$. The second term is the singular local part of §S4: in momentum space it is an algebraic, frequency-dependent constant tensor, and keeping it there is what lets us avoid a real-space principal value entirely. The radiation shell is

$$q_y^2 + q_z^2 + \beta^2 = k_c^2, \quad (\text{S28})$$

which after integration over transverse momentum becomes the branch continuum of §S7. For real $|\beta| > k_c$ —the guided region—the denominator in (S27) never vanishes on the real transverse plane, which is why guided modes are ordinary poles rather than resonances.

S5.3 Cross-section basis and Galerkin matrices

On $|y| < W/2$ use normalized Legendre functions

$$\phi_m^{(W)}(y) = \sqrt{\frac{2m+1}{W}} P_m(2y/W), \quad (\text{S29})$$

whose exact Fourier transform is

$$\Phi_m^{(W)}(q_y) = \sqrt{W(2m+1)} (-i)^m j_m(q_y W/2) \quad (\text{S30})$$

with j_m a spherical Bessel function [8]. Using the analogous basis in z and three Cartesian components, $\phi_{mns} = \phi_m^{(W)}(y)\phi_n^{(T)}(z)\mathbf{e}_s$, which is complete in the vector L^2 of the rectangular cross-section, the cladding Green matrix is the explicit double integral

$$[G_c^{(N)}]_{mns, m'n's'} = \int \frac{dq_y dq_z}{(2\pi)^2} \Phi_m^* \Phi_n^* [\widetilde{\mathbf{G}}_c]_{ss'} \Phi_{m'} \Phi_{n'}. \quad (\text{S31})$$

Every element is a convergent two-dimensional integral of Bessel and rational functions, with no singularity on the real (q_y, q_z) plane once β lies on the deformed contour.

S5.4 Dyson resummation of the core

The beam is restored by resumming its own susceptibility $\chi_b = \epsilon_b - \epsilon_c$ over the rectangular support,

$$\widetilde{\mathbf{G}}_{\text{wg}} = \widetilde{\mathbf{G}}_c + \widetilde{\mathbf{G}}_c k_0^2 \chi_b \widetilde{\mathbf{G}}_{\text{wg}} \implies \widetilde{\mathbf{G}}_{\text{wg}} = [I - k_0^2 \chi_b \widetilde{\mathbf{G}}_c]^{-1} \widetilde{\mathbf{G}}_c. \quad (\text{S32})$$

The finite sequence is $\mathbf{G}_{\text{wg}}^{(N)} = [I - k_0^2 \chi_b G_c^{(N)}]^{-1} G_c^{(N)}$ with $N = (N_y, N_z, N_q)$, in which N_q controls the transverse-momentum quadrature in (S31). No absorber thickness or strength appears anywhere in (S32).

S5.5 Guided poles are not inserted by hand

The guided propagation constants are the zeros of

$$\det A_{\text{wg}}(\beta, \omega) = 0, \quad A_{\text{wg}} = I - k_0^2 \chi_b \widetilde{\mathbf{G}}_c, \quad (\text{S33})$$

i.e. poles of the same resolvent that carries the radiation continuum. At a simple guided pole β_g with null vectors x_g, y_g ,

$$\text{Res}_{\beta=\beta_g} \widetilde{\mathbf{G}}_{\text{wg}} = \frac{x_g y_g^\dagger \mathbf{G}_c(\beta_g)}{y_g^\dagger \partial_\beta A_{\text{wg}}(\beta_g) x_g}, \quad (\text{S34})$$

so the discrete term of the real-space Green tensor arrives with its normalization fixed by the pole derivative. This matters: open-waveguide modes have no finite-box norm, and normalizing them by one is a standard route to an inconsistent radiation weight [9]. Deforming the contour then splits (S26) into

$$\mathbf{G}_{\text{wg}}^+ = i \sum_{g,+} e^{i\beta_g |x-x'|} \text{Res}_{\beta_g} \widetilde{\mathbf{G}}_{\text{wg}} + G_{\text{rad}}^+ + G_{\text{ev}}, \quad (\text{S35})$$

the three terms being guided propagation, the outgoing radiation continuum, and the evanescent remainder.

S5.6 On-shell evaluation inside the light cone

For real β strictly inside the light cone, $\kappa^2 = k_c^2 - \beta^2 > 0$ and the transverse radial integral in (S31) meets the shell. Writing $H(q)$ for the basis-projected numerator, the identity used is

$$\int_0^Q \frac{q H(q) dq}{q^2 - \kappa^2 - i0} = \int_0^Q \frac{q [H(q) - H(\kappa)]}{q^2 - \kappa^2} dq + \frac{H(\kappa)}{2} \log \frac{Q^2 - \kappa^2}{\kappa^2} + \frac{i\pi}{2} H(\kappa). \quad (\text{S36})$$

The first integrand is regular at $q = \kappa$; the last term is the exact on-shell radiation contribution and changes sign for the incoming boundary value. No small artificial loss is added anywhere. The radiation spectral density

$$\rho_{\text{rad}}(\beta) = \frac{\widetilde{\mathbf{G}}_{\text{wg}}^+(\beta) - \widetilde{\mathbf{G}}_{\text{wg}}^-(\beta)}{2\pi i} \quad (\text{S37})$$

must be Hermitian positive semidefinite for a passive real-frequency waveguide, and this is an internal test rather than an assumption.

S6 Reference Green tensor: stratified stack

S6.1 Why the homogeneous reference fails here

The second device is not a beam in one cladding. From bottom to top it is: an SiO_2 substrate half-space, $n = 1.45$; a 300 nm layer that contains the patterned Si_3N_4 strip, $n = 2.0$, with air elsewhere in that layer; 200 nm anthracene, $n = 1.8$; 200 nm PVA, $n = 1.5$; and an air half-space. The films are invariant in x and y ; only the strip is localized.

Replacing the stack by an effective cladding index deletes the guided and leaky poles of the films. Those poles are precisely the channels through which the cavity loses energy vertically, so an effective-index reference cannot support a Q optimization even if it reproduces the resonance wavelength. The correct reference is therefore the *planar stack*, with the strip treated as a localized susceptibility inside the source layer and restored by the same Dyson step as (S32).

The stack also breaks z reflection. The z -parity sector used for the air-clad device is not closed under the multilayer Maxwell operator. Exact x reflection of the symmetric cavity and y reflection of the cross-section survive, so the symmetry projector of §S9 retains a fixed (x, y) sector while carrying both z parities. This roughly doubles the vertical basis at fixed accuracy and is the main reason the second device is more expensive per pole solve than the first.

S6.2 Generalized Fresnel recursion

At fixed in-plane momentum $k_{\parallel}^2 = \beta^2 + q_y^2$, the vertical wavenumber in layer j is

$$k_{zj} = \sqrt{(n_j k_0)^2 - k_{\parallel}^2}, \quad \text{Im } k_{zj} \geq 0, \quad (\text{S38})$$

the branch being fixed by decay rather than by a principal value. TE and TM Fresnel coefficients are composed recursively through each finite layer: if R_j is the reflection seen beyond a layer of thickness d_j , then

$$R_{j-1} = \frac{r_{j-1,j} + R_j e^{2ik_{zj}d_j}}{1 + r_{j-1,j}R_j e^{2ik_{zj}d_j}}. \quad (\text{S39})$$

The recursion retains the exact film poles—the zeros of the denominator—and is holomorphic in every thickness and index away from the branch points of (S38). This is what makes the whole stack differentiable in principle, even though the campaigns here hold the film stack fixed.

S6.3 Scalar vertical kernels and the Wronskian

Let the source layer occupy $0 < z < d$ and let r_b, r_t be the generalized reflections seen at its bottom and top from inside. The two solutions satisfying the respective outgoing reflection conditions are

$$u_L(z) = e^{-ik_z z} + r_b e^{ik_z z}, \quad u_R(z) = e^{ik_z(z-d)} + r_t e^{-ik_z(z-d)}, \quad (\text{S40})$$

with Wronskian

$$W_r = 2ik_z(e^{-ik_z d} - r_b r_t e^{ik_z d}). \quad (\text{S41})$$

The reciprocal scalar kernel normalized by $(\partial_z^2 + k_z^2)g = -\delta$ is then

$$g(z, z') = -\frac{u_L(z_{<})u_R(z_{>})}{W_r}, \quad (\text{S42})$$

and the guided and leaky modes of the stack are the zeros of (S41). Two independent tests are used on this block: g must be symmetric under $z \leftrightarrow z'$, and its jump in $\partial_z g$ across $z = z'$ must equal -1 to machine precision.

The full vector dyadic is assembled from the TE and TM scalar kernels in the standard way, projected onto the asymmetric cross-section basis in q_y and z , and used exactly where $\mathbf{G}_c$ appears for the homogeneous device. For reference, the planar-film TE pole of this stack has $n_{\text{eff}} \approx 1.55246$ and the Si_3N_4 -strip pole $n_{\text{eff}} \approx 1.76502$; both are tracked separately, because they behave very differently under the contour deformation of §S7.

S6.4 Adding the strip

With $\mathbf{G}_b(\beta)$ the projected outgoing Green matrix of the complete planar background, the finite-width Si_3N_4 strip is a compact Dyson problem in exactly the sense of (S32),

$$D_s(\beta) = I - k_0^2 \Delta \epsilon_{\text{core}} \mathbf{G}_b(\beta), \quad \mathbf{G}_{\text{wg}}(\beta) = D_s(\beta)^{-1} \mathbf{G}_b(\beta). \quad (\text{S43})$$

The physical y -odd block of (S43) contains the E_y -dominant feed mode of the cavity. Its pole in β is not a nuisance: it is the mode the cavity is built out of, and its explicit removal from the longitudinal quadrature is the subject of §S10. Attempting the contour integral of (S26) with that pole left in place is the difference between a usable operator and an unusable one, and this is the single largest structural difference between the two devices in this work.

S7 The outgoing contour and its quadrature

S7.1 Branch topology in scaled momentum

Openness is a statement about which side of the branch cuts and poles the longitudinal integration path passes. Working in the scaled variable

$$u = \frac{\beta}{k_c(\omega)}, \quad k_c = n_c \frac{\omega}{c}, \quad (\text{S44})$$

pins the radiation branch points at $u = \pm 1$ for every frequency, while the guided poles sit near $u = \pm n_{\text{eff}}/n_c$ and move only slowly. The outgoing Sommerfeld path must pass *below* every positive- u singularity and *above* its negative- u partner; this is the statement that a disturbance at $x > x'$ propagates to the right and one at $x < x'$ to the left.

S7.2 The deformation

We use the smooth, closed-form path

$$u(t) = t - ih \tanh(t/d) \exp[-(t/L)^8], \quad -\infty < t < \infty, \quad (\text{S45})$$

with dimensional path $\beta(t, \omega) = k_c(\omega) u(t)$. Three properties are used. The $\tanh$ factor gives the required opposite excursions on the two sides of the origin with a single smooth function. The super-Gaussian factor returns the path to the real axis outside $|t| \sim L$, so that the far tails, where the integrand is exponentially small and purely evanescent, are integrated on the real axis where the transverse quadrature is cheapest. And the whole path is independent of ω in the scaled variable, so *holding $u(t)$ fixed while moving ω* is an explicit analytic-continuation homotopy from the causal upper half plane to the resonance.

The depth h and width d are not free: h must exceed the largest imaginary excursion any singularity makes during the continuation, and d must be small enough that the path has separated from $u = \pm 1$ before the branch points are reached. Both are verified per solve by monitoring the distance from the path to every tracked singularity; a crossing appears as a discontinuity in $\tilde{\omega}$ along the homotopy and is rejected rather than accepted with a changed sheet.

S7.3 Orientation test

The first check applied to any implementation of (S45) is the scalar identity

$$\frac{1}{2\pi} \int_{\mathcal{C}_+} \frac{e^{i\beta x}}{\beta^2 - a^2} d\beta = \frac{ie^{iax}}{2a}, \quad x > 0, \quad (\text{S46})$$

which fixes the orientation and the sheet unambiguously and fails loudly for the incoming choice. It is cheap enough to run at every frequency in a convergence study, and it is the test that first exposed the wrong-sheet error described in §S18.

S7.4 Panelized quadrature

The integrand of (S26) for a cavity of length L_{cav} oscillates as $e^{i\beta L_{\text{cav}}}$, so the required node density grows with cavity length. Two quadrature families are implemented. A tangent map compactifies the whole real axis in one rule and is fast, but it hides the omitted high-momentum tail and its convergence cannot be separated from the polynomial order. A panelized finite-cutoff rule is slower and is what all reported numbers use: panels are laid out with breakpoints clustered at $u = \pm 1$ and at the endpoints of the first guided-mode interval, Gauss–Legendre nodes are clustered towards panel endpoints, and the cutoff β_{max}/k_0 is refined independently of the panel width and the node count.

For the second device the accepted rule uses 812 positive- β nodes, a maximum dimensionless panel width of 0.025, four endpoint-clustered nodes per panel, and $\beta_{\max}/k_0 = 5$; halving the node count to 412 changes the imaginary part of the tracked eigenvalue by 0.6 % and Q by 0.65 % (§S14). Agreement obtained from a single tangent-map order is not an error estimate and is not reported as one.

S7.5 Known failure modes

Three ways of getting (S45) wrong were encountered and are recorded here because each produces a plausible number rather than an obvious crash.

1. *Real-axis quadrature at complex frequency.* Substituting $\text{Im } \omega < 0$ into a rule built on the real β axis silently lands on the wrong sheet. The resulting $\tilde{\omega}$ can even be passive; it is simply not the resonance.
2. *Insufficient depth.* If h is too small, a guided pole crosses the path during the continuation and the residue it should have contributed is lost. The symptom is a Q that jumps discontinuously as a geometry parameter is scanned smoothly.
3. *Unclustered panels.* With breakpoints not placed at $u = \pm 1$, the logarithmic and square-root behaviour at the branch points is integrated by a polynomial rule and converges algebraically instead of spectrally. The symptom is a convergence table that improves like N^{-2} and never settles.

S8 Aperture basis and projection

S8.1 Normalized hole coordinates

Inside hole j introduce

$$x - x_j = r_{x,j} \rho \cos \theta, \quad y - y_j = r_{y,j} \rho \sin \theta, \quad 0 \leq \rho \leq 1, \quad (\text{S47})$$

with $z \in [-T/2, T/2]$ and Jacobian $r_{x,j} r_{y,j} \rho$. The map is affine in the semi-axes, which is what makes the basis functions themselves differentiable in the geometry rather than merely defined on a geometry-dependent domain.

S8.2 Zernike–Legendre vector basis

On the unit disk use the real Zernike functions

$$Z_{nm}^c = R_n^m(\rho) \cos m\theta, \quad Z_{nm}^s = R_n^m(\rho) \sin m\theta, \quad (\text{S48})$$

for $0 \leq m \leq n$ with $n - m$ even, the sine family omitted at $m = 0$. Their squared norms on the unit disk are

$$N_{n0} = \frac{\pi}{n+1}, \quad N_{nm} = \frac{\pi}{2(n+1)} \quad (m > 0), \quad (\text{S49})$$

and the radial polynomial is evaluated in the numerically stable Jacobi form

$$R_n^m(\rho) = (-1)^{(n-m)/2} \rho^m P_{(n-m)/2}^{(m,0)}(1 - 2\rho^2), \quad (\text{S50})$$

rather than by the explicit alternating sum, which loses relative accuracy above $n \sim 12$ [10, 8]. Tensoring with the vertical Legendre functions and the three Cartesian unit vectors gives

$$\phi_{j,nm\sigma\ell s} = \frac{Z_{nm}^\sigma(\rho, \theta)}{\sqrt{r_{x,j} r_{y,j} N_{nm}}} \sqrt{\frac{2\ell+1}{T}} P_\ell(2z/T) \mathbf{e}_s, \quad (\text{S51})$$

which is complete in $L^2(H_j)^3$ as the retained degrees grow. At $(n, m, \ell) = (0, 0, 0)$ it reduces to the constant $1/\sqrt{\pi r_{x,j} r_{y,j} T}$, so the lowest order of this basis is algebraically identical to the constant-hole projection used in the earliest revision of the solver—a useful continuity check between development stages.

S8.3 Exact thickness projection

Because the hole spans the complete beam thickness and (S51) uses the same normalized Legendre functions in z as the cross-section basis of (S29),

$$\int_{-T/2}^{T/2} \phi_{i_z}(z) \phi_\ell(z) dz = \delta_{i_z \ell}, \quad (\text{S52})$$

so the vertical projection is exact and contributes no quadrature error. This is a consequence of the through-hole geometry: a partially etched hole would couple vertical orders and require its own quadrature.

S8.4 Projection into the cross-section basis

At fixed β the projection is

$$C_{a,j\mu}(\beta) = \int_{H_j} d^3r \psi_a(y, z) e^{-i\beta x} h_{j\mu}(x, y, z), \quad (\text{S53})$$

which after (S52) reduces to the disk integral

$$\frac{\sqrt{r_x r_y}}{\sqrt{N_{nm}}} e^{-i\beta x_j} \int_0^1 \rho d\rho \int_0^{2\pi} d\theta \phi_{i_y}(r_y \rho \sin \theta) Z_{nm}^\sigma e^{-i\beta r_x \rho \cos \theta}. \quad (\text{S54})$$

The integrand is entire in β , so Gauss–Legendre radial quadrature together with the periodic trapezoidal rule in θ converges spectrally [11]; the angular rule in particular is exponentially accurate because the θ dependence is analytic and periodic. The phase $e^{-i\beta x_j}$ is factored out before quadrature so that a hole displacement is exactly a multiplicative phase, which is what makes (S15) exact at the level of the assembled matrix and not merely at the level of the continuum model.

For the simplest useful case—the constant mode, $\ell = 0$, $(n, m) = (0, 0)$ —the x integral can be done first at fixed y , giving

$$\int dx e^{-i\beta x} h_{j,s} = \frac{e^{-i\beta x_j}}{\sqrt{V_j}} 2a_j(y) \text{sinc}[\beta a_j(y)] e_s, \quad (\text{S55})$$

with $a_j(y) = r_{x,j} \sqrt{1 - y^2/r_{y,j}^2}$ and $V_j = \pi r_{x,j} r_{y,j} T$; the remaining finite y integral is Gauss–Legendre. This closed form is used as a regression test on the general quadrature.

S8.5 Transpose, not conjugate

The projected outgoing Green matrix is

$$\mathbf{G}_H^+(\omega) = \int_{\mathcal{C}_+} \frac{d\beta}{2\pi} C(-\beta)^T \widetilde{\mathbf{G}_{\text{wg}}(\beta, \omega)} C(\beta), \quad (\text{S56})$$

with the left factor $C(-\beta)^T$ and *not* $C(\beta)^\dagger$. For real basis functions and real frequency the two agree, which is exactly why the substitution is tempting and its consequences delayed: the conjugate is not analytic in ω , so a solver built on it converges perfectly well on the real axis and then continues onto nothing in particular. With the transpose, reciprocity predicts $\mathbf{G}_H^+ = (\mathbf{G}_H^+)^T$, which is checked numerically at every assembly (§S14).

S8.6 Production orders

The air-clad device uses a $P2$ hole basis with a single vertical order against a $y3$ – $z1$ cross-section basis; the multilayer device carries both z parities as required by §S6. Convergence under independent increase of N_{xy} , N_z , and N_β is reported in §S14; the orders quoted here are the production settings at which every reported analytical number was computed, not the largest ever run.

S9 Symmetry reduction and finite-array assembly

S9.1 The assembled operator

Concatenating the per-hole blocks of (S51) into a single index set and inserting (S56) into (S24) gives the operator whose root is the cavity,

$$\mathbf{A}(\omega, \boldsymbol{\theta}) = I - k_0^2(\epsilon_c - \epsilon_b) \mathbf{G}_H^+(\omega, \boldsymbol{\theta}), \quad \Delta\epsilon = \epsilon_c - \epsilon_b < 0. \quad (\text{S57})$$

Its dimension is $3N_h \times N_{\text{Zernike}}(N_{xy}) \times (N_z + 1)$ before symmetry reduction: a few thousand for the devices here, which is why a dense factorization is used throughout and why a pole solve is seconds to minutes rather than hours.

S9.2 Symmetry sectors

Mirror symmetry in x is exact for both devices and is used twice: the positive-half parameterization of (S10) halves the number of free coordinates, and the parity projector halves the operator dimension at fixed accuracy. Transverse (y) parity is likewise exact and selects the polarization sector containing the E_y -dominant mode of interest.

Vertical (z) parity is exact only for the air-clad device. The stack of §S6 breaks it, so the multilayer projector retains a fixed (x, y) sector while carrying both z parities. The practical consequence is that a mode which is cleanly labelled by three parities in the first device carries only two labels in the second, and mode tracking (§S11) correspondingly loses one of its cheapest discriminators.

Symmetry is used as a reduction, never as a filter applied after the fact. A mode that leaves its declared sector during a geometry step is a tracking failure and is reported as one.

S9.3 Finite array, not a periodic cell

The array in (S9) is finite and non-uniform: pitch and both radii vary hole by hole. There is no Bloch phase and no supercell anywhere in the construction, and the mirror region is finite, so its residual transmission is part of the computed Q rather than an assumed boundary condition. This is the structural difference from a periodic guided-mode expansion: the object of interest is the finite device, and the leakage out of its ends is one of the loss channels the operator has to represent, alongside vertical radiation.

The cost of a matrix element between holes i and j is dominated by the contour integral (S56), whose integrand carries the phase $e^{i\beta(x_i - x_j)}$. Distant pairs therefore oscillate faster and need more nodes; this, not the matrix dimension, is what sets the cost of a long cavity and what forced the 812-node panelized rule for the 42-hole device (§S7).

S9.4 Block structure and cell cascades

Ordering the index set by hole gives $\mathbf{A}$ a natural block structure. Schur complements on interior blocks produce cell scattering maps that can be cascaded by Redheffer star products, which is the classical way to treat such an array. We use that structure only as a factorization of an already-assembled $\mathbf{A}$, never as an independent model: a cascade of separately computed cell matrices carries an uncontrolled truncation of the shared radiation continuum, which is one of the specific errors catalogued in §S18. Stated positively, the cell picture is available here as a computational convenience whose approximation error is known, because it is derived from the global operator rather than assumed prior to it.

S10 Feed-pole extraction and residues

S10.1 The problem

For the multilayer device, the strip Dyson step (S43) introduces a simple pole of $\mathbf{G}_{\text{wg}}(\beta)$ at the propagation constant β_f of the E_y -dominant guided feed mode. The contour of §S7 passes around it correctly in principle,

but the integrand of (S56) then contains a first-order pole sitting a short distance from a path that must also resolve interference fringes of scale $1/L_{\text{cav}}$. Numerically this is the worst of both worlds: the quadrature must simultaneously resolve a near-singularity and a long oscillation, and the node count required grows without any corresponding gain in accuracy elsewhere.

The remedy is to remove the pole analytically, integrate only what is smooth, and restore the removed part in closed form.

S10.2 Laurent extraction

Let β_f be a simple zero of $\det D_s$ with normalized right and left null vectors r and l . Then

$$\mathbf{G}_{\text{wg}}(\beta) = \frac{R_+}{\beta - \beta_f} + \text{regular}, \quad R_+ = \frac{r [l^H \mathbf{G}_b(\beta_f)]}{l^H D'_s(\beta_f) r}, \quad (\text{S58})$$

the derivative being taken with respect to physical β , so that the residue denominator carries units of inverse length. At the real closed-background pole the operator is Hermitian and $l = r$; away from it the general biorthogonal form is used. Reciprocity fixes the partner residue at $-\beta_f$,

$$R_- = -R_+^T, \quad (\text{S59})$$

which is imposed as a check, not as a definition: the two residues are computed independently and their disagreement is a diagnostic on the whole stratified chain.

The regularized integrand is

$$\mathbf{G}_{\text{wg}}^{\text{reg}}(\beta) = \mathbf{G}_{\text{wg}}(\beta) - \frac{R_+}{\beta - \beta_f} - \frac{R_-}{\beta + \beta_f}, \quad (\text{S60})$$

which is smooth across the whole path and is what the panelized quadrature of §S7 actually integrates.

A symmetric two-sided Laurent estimate of R_+ —evaluating $(\beta - \beta_f)\mathbf{G}_{\text{wg}}$ at $\beta_f \pm \delta$ and averaging—cancels the linear term of the regular part and therefore converges quadratically in δ , which is how R_+ is obtained to the accuracy needed without differentiating D_s symbolically.

S10.3 Exact real-space restoration

The subtracted Laurent terms are not discarded. Closing the contour analytically returns them as ordered feedthrough propagation in real space,

$$\mathbf{G}_{\text{wg}}^{\text{feed}}(x - x') = i R_+ e^{i\beta_f(x-x')} \Theta(x - x') + i R_- e^{-i\beta_f(x-x')} \Theta(x' - x), \quad (\text{S61})$$

with Θ the Heaviside step and the value $\frac{1}{2}$ taken at coincident nodes. The ordering matters: (S61) says that the guided mode propagates rightward from a source to everything on its right and leftward to everything on its left, and applying an unordered symmetric form instead double-counts the guided channel between every pair of holes. In matrix terms, $\mathbf{G}_H^+$ acquires an exact rank-structured contribution assembled from the ordered hole positions, added to the numerically integrated regular part.

S10.4 Magnitude of the effect

The subtraction is not a small correction to a working scheme; without it the scheme does not work. The finite-hole self term of the guided channel alone changes the guided block of the assembled matrix by about 8%, and the near-pole quadrature error it replaces is larger still and is not controlled by refinement at any node count we could afford. The clearest evidence of its importance is the historical comparison in §S18: on one saved geometry direction, the pre-subtraction operator reported an endpoint ratio of 3.6045899—an apparent $3.6\times$ quality-factor improvement—while independent FDTD reported 1.0007628, and the pole-subtracted operator reported 1.0000101. The correct machinery does not merely improve the number; it changes a false positive into a true negative.

S10.5 What must be differentiated

Because (S60) and (S61) both depend on geometry—the former through $C(\beta)$ and the latter through the hole positions in the ordering and through β_f itself—the reverse-mode pass of §S12 must differentiate both. It differentiates the finite-hole Fourier projections in the smooth real-axis continuum and every ordered real-space strip-pole residue. It does not differentiate the rejected legacy contour operator, which is retained in the record but is not on any computational path used here.

S11 Continuation, pole search, and mode identity

S11.1 Locating the first pole

The search begins on the real frequency axis, where the operator is cheapest and its structure best understood. A singular-value scan of $\mathbf{A}(\omega)$ over a frequency window in the declared symmetry sector locates the local minimum of $\sigma_{\min}$; the corresponding left and right singular vectors then seed a local scalar equation

$$\mathbf{v}^T \mathbf{A}(\omega) \mathbf{u} = 0, \quad (\text{S62})$$

which is solved by a secant or Newton iteration in complex ω with $\mathbf{v}$ and $\mathbf{u}$ updated at each step. Because (S62) is scalar, its root finding is inexpensive relative to the assembly of $\mathbf{A}$, and the cost of a pole solve is dominated by the contour integrals of (S56): the measured production figure is 398.0 s for the air-clad device.

Convergence is declared on three criteria simultaneously: the scalar residual $|\mathbf{v}^T \mathbf{A} \mathbf{u}|$ below tolerance; the operator residual $\|\mathbf{A} \mathbf{u}\|/\|\mathbf{u}\|$ at the solve tolerance, measured 2.29×10^{-10} at the production discretization; and $\text{Im } \tilde{\omega} < 0$. A root that satisfies the first two and violates the third is a continuation failure, not a result.

S11.2 Continuation off the real axis

The frequency is moved from the causal upper half plane to the resonance along a homotopy with the scaled contour $u(t)$ of (S45) held fixed. Two quantities are monitored along the path: the distance from the path to each tracked singularity, and the continuity of $\tilde{\omega}$ itself. A singularity crossing produces a step in $\tilde{\omega}$ of order the residue it carries and is rejected. This is the mechanism by which the wrong-sheet error of §S18 is now caught rather than reported.

S11.3 Mode identity under geometry change

Between optimizer steps the geometry changes and the pole must be re-found. Frequency-nearest tracking is insufficient and is the single most dangerous shortcut available here, because a legitimate small shape step can move ω_r by many of the previous linewidths. Recall from (S6) that $|b|/a = 1/(2Q)$: at $Q = 10^4$ a relative frequency shift of 10^{-4} —trivially achievable with a sub-nanometre displacement—is already two linewidths. What must remain continuous is the mode, not a frequency window centred on the old pole.

Candidates are therefore ranked by, in order:

1. biorthogonal field overlap with the accepted QNM, evaluated on a common representation;
2. localization in the intended cavity region;
3. symmetry and polarization sector;
4. agreement with the first-order complex-pole prediction from (S63);
5. frequency proximity, as a final tie-breaker only.

If no candidate clears the overlap and localization gates, the step is declared *unresolved*. An unresolved step does not update the quasi-Newton history and does not enter any reported trajectory. This rule exists because a silent mode switch makes a correct positive directional derivative appear to produce a negative objective change, which will then poison a curvature update for many subsequent steps.

Two representative overlap figures from the record: a physical-template overlap of 0.92776 between the analytical mode and the independently computed field template for the air-clad baseline, and branch overlaps above 0.999614 across the accepted line search of the multilayer device (§S14).

S11.4 Multiple poles and near-degeneracies

Equation (S63) assumes a simple root. Near a degeneracy the normalization $\mathcal{M}$ of (S5) approaches zero and the derivative diverges, which is physically correct—the modes are exchanging identity—and numerically fatal. The implementation monitors $|\mathcal{M}|$ relative to $\|\partial_\omega \mathbf{A}\| \|\mathbf{v}\| \|\mathbf{u}\|$ and flags any step at which that ratio falls below a fixed threshold. No such flag was raised on either campaign reported here, which is consistent with both devices being single-mode in their declared sector over the wavelength window used, but it is a limitation of the present treatment rather than a general guarantee: a strongly multimode cavity would need the block generalization of (S63) and is outside the scope of this work.

S12 Pole and quality-factor derivatives

S12.1 The implicit derivative

From (S4), at a simple pole with null vectors $\mathbf{v}, \mathbf{u}$,

$$\boxed{\frac{\partial \tilde{\omega}}{\partial \theta_j} = -\frac{\mathbf{v}^T (\partial_{\theta_j} \mathbf{A}) \mathbf{u}}{\mathbf{v}^T (\partial_\omega \mathbf{A}) \mathbf{u}}} \quad (\text{S63})$$

This is exact for the discrete operator: no linearization of the geometry, no perturbative treatment of the radiation, and no assumption that the mode is weakly perturbed. It requires only that the root be simple and that $\partial_{\theta_j} \mathbf{A}$ be available.

Separating real and imaginary parts of $\tilde{\omega} = a + ib$,

$$\frac{\partial Q}{\partial \theta_j} = -\frac{1}{2b} \frac{\partial a}{\partial \theta_j} + \frac{a}{2b^2} \frac{\partial b}{\partial \theta_j}, \quad (\text{S64})$$

and for the objective actually used,

$$\frac{\partial \log Q}{\partial \theta_j} = \frac{1}{a} \frac{\partial a}{\partial \theta_j} - \frac{1}{b} \frac{\partial b}{\partial \theta_j}. \quad (\text{S65})$$

Because $|b|/a = 1/(2Q)$, the second term of (S65) carries the signal and the first is nearly irrelevant: a construction that gets a right to twelve digits and b right to three is useless as a Q optimizer, and this is why the tests of §S14 are stated on the Q derivative rather than on the pole.

S12.2 Assembling $\partial_{\theta_j} \mathbf{A}$

From (S57) the geometry enters $\mathbf{A}$ only through $\mathbf{G}_H^+(\omega, \boldsymbol{\theta})$, and from (S56) that dependence is carried entirely by the projection matrices $C(\pm\beta; \boldsymbol{\theta})$ —plus, for the multilayer device, by the ordered feed-through term (S61). Differentiating (S56),

$$\begin{aligned} \partial_{\theta_j} \mathbf{G}_H^+ = \int_{C_+} \frac{d\beta}{2\pi} \Big[(\partial_{\theta_j} C(-\beta))^T \widetilde{\mathbf{G}_{\text{wg}}} C(\beta) \\ + C(-\beta)^T \widetilde{\mathbf{G}_{\text{wg}}} (\partial_{\theta_j} C(\beta)) \Big] + \partial_{\theta_j} \mathbf{G}_H^{\text{feed}}, \end{aligned} \quad (\text{S66})$$

with the contour and $\widetilde{\mathbf{G}_{\text{wg}}}$ independent of geometry. The projection derivatives come from (S15) and (S17): a centre displacement is a multiplicative $-i\beta$, and a radius derivative is the analytic Bessel expression, both evaluated on exactly the same quadrature nodes as the primal. There is no finite differencing anywhere in the chain and no remeshing.

S12.3 Reverse-mode contraction

Equation (S63) shows that only the scalar $\mathbf{v}^T(\partial_{\theta_j}\mathbf{A})\mathbf{u}$ is ever needed, not the matrix $\partial_{\theta_j}\mathbf{A}$ itself. Writing $\Lambda = \mathbf{u}\mathbf{v}^T$, the required scalar is $\text{Tr}[\Lambda^T\partial_{\theta_j}\mathbf{A}]$, and the entire gradient is one reverse pass that seeds Λ at the output and accumulates through the primitives—contour quadrature, Dyson inverse, disk quadrature, Bessel form factors, pitch-to-centre accumulation—in reverse order. The cost is a small constant multiple of one assembly and is independent of the number of geometry coordinates.

Two implementation points are worth recording. First, the pitch parameterization (S10) means the adjoint of the centre-accumulation step is a reverse cumulative sum, so the derivative with respect to an inner pitch automatically collects the contributions of every hole outboard of it. Second, the ordered Heaviside structure of (S61) is differentiated as written, with the $\frac{1}{2}$ at coincident nodes, since the ordering is fixed by hole index and does not change under the small displacements a line search applies.

Measured cost: 370.0s for the full 45-coordinate gradient of the air-clad device, against 398.0s for the pole solve it differentiates.

S12.4 Boundary form

Equation (S63) can be rewritten as a boundary integral, which is useful both as an independent implementation and as a conceptual statement about which fields matter. For an interface Γ separating isotropic media 1 and 2, with unit normal $\mathbf{n}$ from 1 to 2 and a positive scalar velocity h moving Γ along $\mathbf{n}$,

$$\mathbf{v}^T(\partial_p\mathbf{A})\mathbf{u} \propto \int_{\Gamma} h \left[(\epsilon_1 - \epsilon_2) \mathbf{E}_{L\parallel} \cdot \mathbf{E}_{R\parallel} - (\epsilon_1^{-1} - \epsilon_2^{-1}) D_{L\perp} D_{R\perp} \right] dA, \quad (\text{S67})$$

tangential $\mathbf{E}$ and normal $\mathbf{D}$ appearing because those are the components continuous across a sharp dielectric interface [12, 13]. Magnetic contrast contributes the dual form in $\mathbf{H}$ and $\mathbf{B}$ and is absent here. Writing the bracket as a scalar shape-gradient density $g_{\omega}(s)$,

$$\delta\tilde{\omega} = \int_{\Gamma} g_{\omega}(s) h(s) dA \quad (\text{S68})$$

for any admissible normal velocity—which is why hole radii, hole positions, and beam sidewalls are all differentiable in one framework: each is a piece of the same movable boundary.

The proportionality constant and the overall sign in (S67) depend on the precise definition of $\mathbf{A}$, on the orientation of $\mathbf{n}$, and on whether p is a displacement or a signed-distance value. We fix them by test rather than by argument:

1. translate one planar interface by a known signed amount;
2. compare the predicted complex-frequency shift against a centred difference on identical connectivity;
3. reverse the normal and verify that the gradient sign reverses;
4. repeat for every material pair in the device.

S12.5 Boundary metric

The raw density g_Q is generally too rough to advect a boundary directly. It is mapped through a Sobolev metric

$$(I - \ell^2 \partial_s^2)h = g_Q, \quad (\text{S69})$$

with s arclength and ℓ a fabrication-aware smoothing length. This is a change of inner product, not an additional objective term: the ascent direction it produces is still an ascent direction for the same Q . Separate hole and sidewall blocks may use different ℓ , but their RMS normalizations must be declared explicitly so that one block does not dominate merely by carrying more sample points.

The geometry handed to the next solve must be the geometry the motion actually produced. Smoothing, signed-distance reinitialization, polygon repair, and remeshing can each introduce a displacement that does not vanish as the requested step does, and a line search is only meaningful when

$$\text{dist}(\Gamma(\alpha), \Gamma(0)) \rightarrow 0 \quad \text{as} \quad \alpha \rightarrow 0. \quad (\text{S70})$$

For the parameterization of §S3 this holds by construction. It did not hold for an earlier mesh-based revision, which is documented in §S18.

S12.6 Trust ladder

Not every derivative in a project earns the same claim. We use an explicit five-level ladder and state which level each reported quantity has reached.

level	evidence	permitted claim
algebraic	adjoint agrees with a discrete directional derivative	the implementation differentiates its discrete operator
local surrogate	predicted and solved pole changes agree on common connectivity at high mode overlap	the surrogate has a valid local ascent direction
representation transfer	the saved contour reproduces on an independent mesh	the boundary handoff is not dominating the step
physical transfer	periodic independent-solver evaluation improves under a fixed protocol	the optimization improves a physical observable
converged device	Q stable to mesh, domain, run time, and fit choices	a numerical Q may be quoted as validated

In this work the Q shape derivative has passed the first four levels; the fifth is exactly what the panel of §S15 exists to establish and is why both device gains remain provisional. Mode volume and coupling derivatives have *not* passed their own algebraic and local tests to the same standard and are therefore excluded from the quasi-Newton objective entirely; V_{eff} appears in this work only as a reported diagnostic, never as an optimized quantity.

S13 Optimizer, trust region, and acceptance gates

S13.1 Objective and variables

The optimizer maximizes $\log Q$ over the geometry vector $\boldsymbol{\theta}$ of §S3, with the beam width and the layer thicknesses held fixed. Bound constraints keep every semi-axis positive and prevent adjacent holes from overlapping; they were not active at any accepted step in either campaign.

Quasi-Newton curvature is accumulated by BFGS (air-clad device: a short constrained ascent; multilayer device: a pure BFGS run of 105 accepted steps out of 106 attempted). Curvature is updated only after a clean accepted secant—an unresolved or rejected step contributes nothing—because a single poisoned pair persists in the approximate inverse Hessian for many subsequent iterations.

S13.2 Step control

Each proposed step is scaled so that its largest single-coordinate displacement does not exceed a declared bound. For the air-clad campaign the bound was 1 nm of scaled L^2 displacement per step, which realized a maximum single-coordinate move of about 0.34 nm. This is deliberately conservative: it keeps the step inside the regime where the first-order prediction from (S63) is checkable against the solved pole, which is the local-surrogate level of the trust ladder.

A line search evaluates candidate poles on common connectivity—the same quadrature nodes, the same basis, the same contour—so that the measured change is a property of the geometry and not of a discretization. A representative accepted line search from the multilayer device:

step	max coord. change	predicted Q_{ana}	branch overlap
−0.50 nm	−0.50 nm	5 492.77	0.999614
−0.25 nm	−0.25 nm	6 348.88	0.999905
baseline	0	7 270.34	1.000000
+0.25 nm	+0.25 nm	8 189.34	0.999907
+0.50 nm	+0.50 nm	8 997.81	0.999632

A +1 nm probe predicted a still larger Q but its branch overlap fell below the gate of 0.999 and was excluded. That exclusion is the point of the gate: the larger step was not obviously wrong, it simply could no longer be certified to be the same mode.

S13.3 Acceptance gates

A step is accepted only if all of the following hold.

1. *Pole certified.* Passive sign, residual at tolerance, complex symmetry within the running bound (§S11).
2. *Mode resolved.* Biorthogonal overlap and localization gates cleared, declared symmetry sector retained, branch overlap above 0.999.
3. *Geometry valid.* All bounds satisfied; no hole overlap; the applied displacement equals the requested displacement, i.e. (S70).
4. *Surrogate gain positive.* The solved $\log Q$ at the applied geometry exceeds the previous accepted value. A merely predicted gain is not sufficient.

Failure of any gate makes the step *unresolved* rather than merely rejected: the trust radius shrinks, no curvature is recorded, and the state is logged. Of 106 attempted multilayer iterations, 105 were accepted.

S13.4 One loop, written out

For a single accepted geometry the loop is:

1. solve and certify the QNM (§S11);
2. compute the complex boundary gradient from matched left/right fields (§S12);
3. regularize it in the declared boundary metric (S69);
4. predict $\delta\tilde{\omega}$ and $\delta\log Q$ using the *applied* boundary motion, not the requested one;
5. solve candidate poles on common connectivity for the line search;
6. reject unresolved tracking, invalid geometry, poor transfer, or a nonpositive measured surrogate gain;
7. update the quasi-Newton history only after a clean accepted secant;
8. periodically rebuild the geometry independently and evaluate it with an independent solver (§S15).

Steps 5 and 8 answer different questions and are never substituted for one another. Common connectivity makes the local derivative *measurable*; independent reconstruction tests whether the representation and its discretization *transfer*.

S13.5 What the independent solver is not allowed to do

The independent FDTD layer is a qualification instrument. It supplies no gradient, participates in no line search, and cannot cause a step to be accepted. Its diagnostics are recorded at a fixed cadence—every three accepted steps in the multilayer campaign—and are read after the fact. This separation is what makes the transfer statement of the main text meaningful: if FDTD results had been allowed to influence step selection, agreement between the two would be partly circular.

Evidence

S14 Analytical convergence and gradient validation

This section reports what the analytical construction has been shown to do. It establishes the first two levels of the trust ladder of §S12 and no more: nothing here bears on whether the truncated operator’s absolute Q equals a converged Maxwell solution.

S14.1 Operator identities

At every accepted state the following are computed and recorded.

Passivity.

$\text{Im } \tilde{\omega} < 0$. A violation is a continuation failure and aborts the step.

Residual.

$\|\mathbf{A}\mathbf{u}\|/\|\mathbf{u}\|$ at the solve tolerance; measured 2.29×10^{-10} for the air-clad production operator.

Null-vector equations.

Maximum left and right equation errors below 2.35×10^{-15} .

Complex symmetry.

$\|\mathbf{A} - \mathbf{A}^T\|/\|\mathbf{A}\|$, predicted zero by reciprocity; measured between 1.1×10^{-16} and 3.8×10^{-16} across all accepted multilayer runs, and 3.05×10^{-14} for the air-clad preflight.

Residue reciprocity.

The scattering-residue partner relation (S59), measured 1.53×10^{-15} .

Contour orientation.

The closed-form identity (S46), which is exact and therefore checked to machine precision.

The air-clad release preflight, run on the packaged production operator, gave a passive pole $\tilde{\omega}/\omega_0 = 0.9850125798555 + -6.3776704224 \times 10^{-5} \text{ i}$, $Q_{\text{ana}} = 7\,722.35$, and a physical-template overlap of 0.92776, with the full 45-coordinate reverse-mode gradient completing in 370.0 s after a 398.0 s pole solve.

S14.2 Refinement convergence

The three truncation indices of (S25) are refined independently, because agreement obtained by moving two at once is not an error estimate. For the multilayer device the accepted longitudinal rule uses 812 positive- β nodes with maximum dimensionless panel width 0.025, four endpoint-clustered nodes per panel, and $\beta_{\text{max}}/k_0 = 5$. Halving the node count and doubling the panel width gives:

quantity	coarse	accepted	change
max panel width	0.05	0.025	2× refinement
β nodes	412	812	1.97×
Im tracked eigenvalue	0.0018993	0.00188793	0.6 %
Q_{ana} (directional-gradient run)	7 223.48	7 270.74	0.65 %
cutoff β_{max}/k_0	4	5	—

A sub-percent change in the imaginary part under a factor-two refinement of the quantity that generates it is the strongest available statement about this block. It is not a claim of absolute accuracy: the ratios in §S17 show the truncated operator disagreeing with an independent solver by more than an order of magnitude at some geometries, and the two statements are compatible because they concern different things—the internal convergence of a truncation, and its distance from physical reality.

S14.3 Gradient validation

The analytic reverse-mode gradient is compared against central finite differences taken on identical connectivity, with the geometry direction taken from a saved optimizer step and scaled to a 0.1 nm maximum

coordinate displacement. For the multilayer device at the development grid, the tracked eigenvalue derivative gave

central finite difference	$-9.9476613141 \times 10^{-5} + 6.169167943 \times 10^{-8} i$
analytic VJP	$-9.9476626041 \times 10^{-5} + 6.1700312856 \times 10^{-8} i$

and the corresponding Q derivatives were 3.0939549116 by central difference against 3.0939341486 by the analytic vector–Jacobian product, a relative error of 6.71×10^{-6} .

The Q derivative is the meaningful test, not the eigenvalue derivative: the imaginary parts above agree to about four digits while the real parts agree to eight, and (S65) weights the former far more heavily. Agreement at 6.71×10^{-6} on the Q derivative is therefore a statement about the part of the calculation that actually drives the optimizer.

S14.4 A directional gate

Beyond a pointwise derivative check, the saved iteration-0-to-30 geometry direction was magnified until its largest coordinate changed by 1 nm and scanned:

direction coordinate	Q_{ana}	branch overlap
−1	7 217.44	0.99999942
0	7 223.48	1.00000000
+1	7 226.87	0.99999943

The actual saved iteration-30 displacement is 1/63.39 of that test step, so the rescaled predicted endpoint ratio is ≈ 1.0000103 . That is consistent with the direct checkpoint calculation and, importantly, correctly classifies the direction as flat rather than as an improvement—see §S18, where the pre-repair operator classified the same direction as a $3.6\times$ gain.

S14.5 Complete audited value table

Every value quoted in this section is extracted from the retained gate records by a script that verifies the literal string still appears in the cited file, and fails otherwise.

quantity	value	source record
quan_operator_residual	2.29e-10	VALIDATION_EVIDENCE.md
quan_reciprocity_error	3.05e-14	VALIDATION_EVIDENCE.md
quan_template_overlap	0.92776	VALIDATION_EVIDENCE.md
quan_null_vector_error	2.35e-15	VALIDATION_EVIDENCE.md
quan_residue_reciprocity	1.53e-15	VALIDATION_EVIDENCE.md
quan_pole_solve_seconds	398.0 s	VALIDATION_EVIDENCE.md
quan_gradient_seconds	370.0 s	VALIDATION_EVIDENCE.md
quan_preflight_analytical_q	7722.35	VALIDATION_EVIDENCE.md
quan_unverified_continuation_q	30307.26	VALIDATION_EVIDENCE.md
d23_beta_nodes_coarse	412	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_beta_nodes_accepted	812	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_panel_width	0.025	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_beta_cutoff	5 k0	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_im_eigenvalue_coarse	0.0018993	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_im_eigenvalue_accepted	0.00188793	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_im_eigenvalue_change_percent	0.6 %	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_q_coarse	7223.48	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_q_accepted	7270.74	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_q_change_percent	0.65 %	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_reciprocity_min	1.1e-16	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_reciprocity_max	3.8e-16	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_gate_fdt_d_q	6176.55	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_gate_analytical_q	7270.34	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_grad_relative_error	6.71e-06	REPAIRED_OPTIMIZER_FDTD_GATE.md
d23_grad_fd_q_derivative	3.0939549116	REPAIRED_OPTIMIZER_FDTD_GATE.md
d23_grad_vjp_q_derivative	3.0939341486	REPAIRED_OPTIMIZER_FDTD_GATE.md

quantity	value	source record
d23_flat_direction_fdt_ratio	1.0007628	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_flat_direction_legacy_ratio	3.6045899	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_flat_direction_repaired_ratio	1.0000101	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_line_search_min_overlap	0.999614	REPAIRED_OPTIMIZER_FDTD_GATE.md
d23_branch_overlap_gate	0.999	REPAIRED_OPTIMIZER_FDTD_GATE.md
quan_pole_real	0.9850125798555 w0	VALIDATION_EVIDENCE.md
d23_n_substrate	1.45	DERIVATION_PHASE1.md
d23_n_strip	2.0	DERIVATION_PHASE1.md
d23_n_anthracene	1.8	DERIVATION_PHASE1.md
d23_n_pva	1.5	DERIVATION_PHASE1.md
d23_t_anthracene	200 nm	DERIVATION_PHASE1.md
d23_t_pva	200 nm	DERIVATION_PHASE1.md
d23_t_source_layer	300 nm	DERIVATION_PHASE1.md
d23_gate_pole_real	0.9992524494 w0	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_gate_pole_imag	-6.871739e-05 w0	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_gate_lambda_analytical	796.28805 nm	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_gate_lambda_fdt	796.6645 nm	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_gate_q_error_percent	17.72 %	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_gate_lambda_error_percent	-0.0473 %	SUBTRACTED_REAL_AXIS_QNM_GATE.md
quan_phase57_analytical_q	14568.9	VALIDATION_EVIDENCE.md
quan_phase57_fdt_baseline	15468.56	VALIDATION_EVIDENCE.md
quan_phase57_fdt_optimized	42792.17	VALIDATION_EVIDENCE.md
quan_phase57_ratio	2.7664	VALIDATION_EVIDENCE.md
quan_phase58_step1_q	15662.82	VALIDATION_EVIDENCE.md
quan_phase58_step2_q	16818.3	VALIDATION_EVIDENCE.md
quan_phase58_max_coord_nm	0.34 nm	VALIDATION_EVIDENCE.md
quan_pole_imag	-6.3776704224e-05 w0	VALIDATION_EVIDENCE.md
d23_panel_width_coarse	0.05	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_node_refinement_factor	1.97 x	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_beta_cutoff_coarse	4 k0	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_scan_q_minus	7217.44	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_scan_q_zero	7223.48	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_scan_q_plus	7226.87	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_scan_overlap_minus	0.99999942	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_scan_overlap_plus	0.99999943	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_scan_rescale	63.39	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_rescaled_endpoint_ratio	1.0000103	SUBTRACTED_REAL_AXIS_QNM_GATE.md
d23_fd_eigen_real	-9.9476613141e-05	REPAIRED_OPTIMIZER_FDTD_GATE.md
d23_fd_eigen_imag	6.169167943e-08	REPAIRED_OPTIMIZER_FDTD_GATE.md
d23_vjp_eigen_real	-9.9476626041e-05	REPAIRED_OPTIMIZER_FDTD_GATE.md
d23_vjp_eigen_imag	6.1700312856e-08	REPAIRED_OPTIMIZER_FDTD_GATE.md
d23_ls_q_minus_half	5492.77	REPAIRED_OPTIMIZER_FDTD_GATE.md
d23_ls_q_minus_quarter	6348.88	REPAIRED_OPTIMIZER_FDTD_GATE.md
d23_ls_q_baseline	7270.34	REPAIRED_OPTIMIZER_FDTD_GATE.md
d23_ls_q_plus_quarter	8189.34	REPAIRED_OPTIMIZER_FDTD_GATE.md
d23_ls_q_plus_half	8997.81	REPAIRED_OPTIMIZER_FDTD_GATE.md
d23_ls_overlap_minus_half	0.999614	REPAIRED_OPTIMIZER_FDTD_GATE.md
d23_ls_overlap_minus_quarter	0.999905	REPAIRED_OPTIMIZER_FDTD_GATE.md
d23_ls_overlap_plus_quarter	0.999907	REPAIRED_OPTIMIZER_FDTD_GATE.md
d23_ls_overlap_plus_half	0.999632	REPAIRED_OPTIMIZER_FDTD_GATE.md

S15 Independent-FDTD protocol and the ringdown audit

S15.1 Role and separation

The independent solver is a finite-difference time-domain code with perfectly-matched-layer absorbers [14, 15, 16]. It is used only to evaluate frozen geometries. It supplies no gradient, participates in no line search, and no analytical step was ever accepted or rejected on the basis of an FDTD result (§S13). This separation is what makes any observed agreement informative rather than circular.

S15.2 Endpoint procedure

For a comparison between two geometries, both are rendered from the retained geometry vectors through the same generator into simulations that are identical in source spectrum and position, monitor set and

position, domain extent, mesh density, absorber configuration, run time, and resonance-analysis code. Only the geometry differs. Quality factors come from an exponential fit to the time-monitor ringdown after the source has decayed, together with a reported fit error; effective mode volume and a localization fraction are recorded alongside.

S15.3 Ringdown adequacy

A ringdown fit constrains the decay rate only if the recorded window is long compared with the amplitude lifetime of the mode. With

$$\gamma = \frac{\omega}{2Q}, \quad \tau = \frac{1}{\gamma} = \frac{2Q}{\omega}, \quad \mathcal{N} = \frac{t_{\text{run}} - t_{\text{start}}}{\tau}, \quad (\text{S71})$$

$\mathcal{N}$ counts the lifetimes actually observed. We audited all 35 retained FDTD observations across both campaigns; the complete table appears at the end of this section. 24 of them have $\mathcal{N} < 1$.

Three remarks are needed for this to be read correctly.

1. The air-clad records store a fitted decay rate directly. It reproduces $\omega/(2Q)$ to machine precision, so the two are the same fitted quantity in two forms; the agreement is bookkeeping, not an independent check of the fit.
2. The multilayer diagnostics do not retain a source start time. Their window is therefore bounded above by the total run time and the reported $\mathcal{N}$ is optimistic—the true value is smaller by the source ramp, which for comparable settings on the other device was 1.35 ps.
3. $\mathcal{N} < 1$ does not make a fit meaningless; it makes it poorly conditioned and, in this data set, biased towards larger Q .

The internal evidence for that bias direction is direct. Along the multilayer trajectory, the 12 ps diagnostic at step 61 reports $Q_{\text{FDTD}} = 156\,520$; every subsequent 64 ps analysis, at geometries with strictly higher analytical Q , reports values near 58 240.27. Either the optimizer destroyed a factor of three in physical Q over thirteen steps while its own objective rose, or the short-window fits were overestimating. The protocol change was made at the time for exactly this reason.

S15.4 Predeclared promotion panel

Promotion of a device gain from provisional to validated requires a matched convergence panel that was declared before execution and recorded in machine-readable form. Each endpoint is evaluated over the full factorial of four binary axes:

axis	air-clad	multilayer
steps per wavelength	28 / 35	18 / 20
run time (ps)	60 / 90	64 / 96
PML layers	12 / 16	12 / 16
padding	1.25 / 1.5625	1.0 / 1.25

That is 16 variants per endpoint and 64 tasks in total, with nine required artifacts per task: the immutable simulation description, solver version and task identifier, realized cost, the resonance table with fit diagnostics, mode identity and localization, a geometry plot, a field plot carrying its quantity, slice, normalization and colour scale, the spectrum and ringdown, and artifact checksums.

The promotion gate requires all of: matched endpoint settings; stable mode identity across every variant; passing fit diagnostics; the optimized Q above the initial Q in *every* executed variant; and a central value selected by the predeclared refinement rule. The panel status is predeclared not submitted, and no build command in the source repository can submit it.

S15.5 A required amendment to the panel

The adequacy audit revises the declaration. At the air-clad optimized endpoint’s amplitude lifetime of 94.5 ps, the predeclared run times of 60 ps and 90 ps give $\mathcal{N} = 0.63$ and 0.95—still under one lifetime. The declared axis cannot certify a Q of that magnitude. The promotion run must extend the run-time axis until $\mathcal{N} \gtrsim 3$ at the optimized endpoint, which at the historical value implies a window beyond 284 ps. For the multilayer endpoint, whose lifetime is 47.65 ps, the declared 64 ps and 96 ps give $\mathcal{N} = 1.3$ and 2, which is better but still short of that target.

We record this as an amendment rather than reinterpreting the original declaration, because the value of a predeclaration lies entirely in its being harder to change than to satisfy.

S15.6 Complete adequacy table

device	record	Q_{FDTD}	$1/\gamma$ (ps)	window (ps)	$\mathcal{N}$	< 1
quan_loncar	quan_baseline	15 449	24.64	22.65	0.919	•
quan_loncar	quan_optimized_phase58	59 337	94.50	22.65	0.240	•
design23	step_0000	18 992	15.98	12.00	0.751	•
design23	step_0003	18 591	15.59	12.00	0.769	•
design23	step_0006	14 900	12.46	12.00	0.963	•
design23	step_0009	14 307	11.93	12.00	1.006	
design23	step_0012	18 699	15.58	12.00	0.770	•
design23	step_0015	19 826	16.53	12.00	0.726	•
design23	step_0018	20 429	17.02	12.00	0.705	•
design23	step_0021	20 411	17.01	12.00	0.705	•
design23	step_0024	20 458	17.07	12.00	0.703	•
design23	step_0027	22 690	18.94	12.00	0.634	•
design23	step_0030	24 151	20.16	12.00	0.595	•
design23	step_0034	22 262	18.59	12.00	0.646	•
design23	step_0037	17 488	14.57	12.00	0.824	•
design23	step_0040	22 573	18.75	12.00	0.640	•
design23	step_0043	39 625	32.91	12.00	0.365	•
design23	step_0046	55 959	46.46	12.00	0.258	•
design23	step_0049	62 126	51.62	12.00	0.232	•
design23	step_0052	66 985	55.61	12.00	0.216	•
design23	step_0055	58 168	48.29	12.00	0.248	•
design23	step_0058	1.0806×10^5	89.77	12.00	0.134	•
design23	step_0061	1.5652×10^5	130.00	12.00	0.092	•
design23	step_0064	1.3578×10^5	112.73	12.00	0.106	•
design23	step_0067	1.4215×10^5	117.92	12.00	0.102	•
design23	step_0070	436.34	0.38	12.00	31.931	
design23	step_0080	54 658	44.87	64.00	1.426	
design23	step_0083	57 106	46.87	64.00	1.365	
design23	step_0086	56 532	46.39	64.00	1.380	
design23	step_0089	57 304	47.02	64.00	1.361	
design23	step_0092	58 481	47.94	64.00	1.335	
design23	step_0096	54 257	44.46	64.00	1.439	
design23	step_0099	55 039	45.08	64.00	1.420	
design23	step_0102	64 176	52.54	64.00	1.218	
design23	step_0105	58 240	47.65	64.00	1.343	

S16 Air-clad Quan–Lončar campaign

S16.1 Device and parameterization

A mirror-symmetric silicon nanobeam, cross-section $700 \text{ nm} \times 220 \text{ nm}$, core index 3.46, suspended in air, with a 990 nm unpatterned lead at each end and 15 elliptical through-holes per side. The free vector is the 15 positive-side pitches and the two semi-axes of each hole: 45 coordinates. The baseline is the deterministic design of Quan and Lončar [17], with a constant 330 nm pitch and radii tapering from 121.27 nm at the cavity centre to 85.75 nm at the mirror; the target wavelength is near 1 502.32 nm in the E_y -dominant sector.

The production analytical model uses the homogeneous air-clad reference of §S5 with a $P2$ hole basis at one vertical order against a $y3$ – $z1$ cross-section basis, and the custom reverse-mode gradient of §S12.

S16.2 A lineage caution

The analytical solver for this device was developed against a *different* benchmark: the encapsulated silicon-nitride nanobeam of Fryett *et al.* [18], with $W = 450$ nm, $T = 330$ nm, $n_b = 2.0$, $n_c = 1.47$, mirror pitch $a_m = 233$ nm, and a reported $Q \simeq 10^5$ near 740 nm. The derivation record therefore quotes those material and geometric values throughout its early phases. They are *not* the production device. Any number in this document that describes the campaign refers to the air-clad silicon device above; where a derivation is cited for its mathematics, the mathematics is device-independent and the constants are not carried over.

That benchmark also contains a geometry ambiguity worth recording, since it affects anyone attempting the same reconstruction: the source describes a linear taper of major diameter and period, mirror diameters of 300 nm/100 nm, an innermost major diameter of 100 nm with central holes “separated by 140 nm”, and a Bragg region of 40 holes. Those statements do not determine a unique geometry—in particular whether 140 nm is a separation or an optimized transverse diameter, and whether 40 is per side or total. The derivation treats the geometry as a parameter vector and retains the earlier reconstruction only as a named candidate.

S16.3 Optimization

The accepted design is step 3 of a constrained ascent, each step a scaled L^2 displacement of at most 1 nm with a realized maximum single-coordinate move of about 0.34 nm. The analytical quality factor along those steps was $14\,568.9 \rightarrow 15\,662.82 \rightarrow 16\,818.3 \rightarrow 18\,037.3$.

A further eight-step continuation of the same ascent reached $Q_{\text{ana}} = 30\,307.26$ but was never evaluated in FDTD. It is retained in the record under a distinct name so that it cannot be confused with the qualified step-3 design, and it supports no claim here: it is evidence of analytical continuation, not of physical improvement.

S16.4 Geometry outcome

The endpoint geometries reduce to the following descriptors, computed from the retained seed files.

quantity	before		after	
	mean	range	mean	range
pitch (nm)	330	330–330	331.59	328–333.7
r_x (nm)	109.74	85.75–121.3	108.95	86.02–122.9
r_y (nm)	109.74	85.75–121.3	109.87	86.57–121.9
r_y/r_x	1	1–1	1.0086	0.9916–1.016
$2r_y/W$	0.3135	0.245–0.3465	0.3139	0.2473–0.3482
$1 - 2r_x/p$	0.3349	0.265–0.4803	0.3427	0.252–0.4796
area (nm ²)	38 225	2.31e+04–4.62e+04	37 986	2.339e+04–4.705e+04

The reading is unambiguous. The mean in-plane aspect ratio moves from 1 to 1.0086; the largest transverse semi-axis change over all 15 holes is 0.819 nm; the mean hole area changes by a factor 0.9937; and the transverse air fraction is essentially unchanged. The only substantial motion is longitudinal, up to 24.6 nm, outward at the mirror end and inward near the centre, i.e. a mild stretch of the taper envelope. This is local refinement of an existing design, and the value of including it here is precisely that it is *not* a dramatic restructuring: it establishes that the method also works in the regime where the answer is already nearly optimal.

S16.5 Independent evaluation

Matched 24 ps FDTD analyses of the two frozen geometries give:

	baseline	Phase-58 step 3
Q_{FDTD}	15 449.47	59 337.16
λ (nm)	1 502.32	1 499.97
Q_{ana}	7 455.16	18 037.3
$Q_{\text{ana}}/Q_{\text{FDTD}}$	0.483	0.304
$\tau = 1/\gamma$ (ps)	24.64	94.5
window (ps)	22.65	22.65
$\mathcal{N}$	0.919	0.24

The FDTD gain is $3.8407\times$ and the analytical gain is $2.42\times$: the surrogate under-predicts both the absolute value and the improvement. A slightly earlier revision of the same comparison, at 12 ps rather than 24 ps, gave $15\,468.56 \rightarrow 42\,792.17$, a ratio of 2.7664; the two ratios differ by 39 % under nothing but a change of ringdown length, which is itself a useful measure of how much the quoted number depends on the protocol.

S16.6 Status

PROVISIONAL. The ratio is a matched comparison, but $\mathcal{N} = 0.24$ at the optimized endpoint means the absolute values are not converged, and the predeclared panel of §S15 has not been executed. The amendment recorded there—extending the run-time axis until $\mathcal{N} \gtrsim 3$ —applies to this device in particular.

S17 Multilayer Design23 campaign

S17.1 Device and parameterization

The stack, bottom to top: SiO_2 substrate half-space, $n = 1.45$; a 300 nm layer containing the patterned Si_3N_4 strip, $n = 2.0$, air elsewhere in that layer; 200 nm anthracene, $n = 1.8$; 200 nm PVA, $n = 1.5$; air half-space. The strip is 556.264 nm wide and carries 42 elliptical through-holes per side, giving 126 free coordinates: 42 pitches and two semi-axes each. The strip width and the layer thicknesses were held fixed. The resonance is near 792.347 nm in the E_y -dominant sector.

The analytical model uses the stratified reference of §S6, the strip Dyson step (S43), the feed-pole subtraction of §S10, and the 812-node panelized contour rule of §S7. Both z parities are carried because the stack breaks vertical reflection.

A superseded revision of this device used 25 holes per side and 75 coordinates; several derivation phases in the retained record describe that version. The production device is the 42-hole, 126-coordinate one, and all numbers here refer to it.

S17.2 Trajectory

A pure BFGS ascent on $\log Q$ accepted 105 of 106 attempted steps, raising Q_{ana} from 8 958.74 to 717 003, a factor of 80.03. Independent FDTD diagnostics were taken every three accepted steps. The complete retained record follows; $\mathcal{N}$ is the ringdown adequacy of (S71) and † marks the single observation excluded by the predeclared fit-quality rule.

step	Q_{ana}	Q_{FDTD}	λ (nm)	V_{eff} $(\lambda/n)^3$	loc.	fit err.	t_{run} (ps)	$\mathcal{N}$
0	8 958.7	18 992	792.35	5.047	0.679	0.000429	12	0.75
3	12 488	18 591	790.03	4.736	0.691	0.000663	12	0.77
6	14 350	14 900	787.81	4.424	0.705	0.00121	12	0.96
9	15 027	14 307	785.35	4.264	0.700	0.000755	12	1.01
12	16 012	18 699	784.56	4.208	0.693	0.00153	12	0.77
15	17 410	19 826	785.26	4.114	0.695	0.00145	12	0.73
18	18 439	20 429	784.56	4.384	0.687	0.00136	12	0.71
21	19 742	20 411	785.09	4.992	0.688	0.00134	12	0.71
24	21 154	20 458	785.67	4.023	0.690	0.00133	12	0.70
27	22 698	22 690	786.12	4.164	0.689	0.000539	12	0.63
30	24 371	24 151	786.35	4.008	0.687	0.00175	12	0.60
34	30 774	22 262	786.29	4.008	0.695	0.000282	12	0.65
37	36 042	17 488	784.78	3.439	0.690	0.00167	12	0.82
40	49 688	22 573	782.13	3.275	0.690	0.00233	12	0.64
43	63 495	39 625	782.25	3.199	0.691	0.000822	12	0.36
46	85 898	55 959	781.94	3.032	0.686	0.00181	12	0.26
49	1.2595×10^5	62 126	782.55	3.412	0.670	0.000222	12	0.23
52	1.5413×10^5	66 985	781.92	3.133	0.668	0.00113	12	0.22
55	1.7445×10^5	58 168	781.92	3.234	0.669	0.00104	12	0.25
58	2.1063×10^5	1.0806×10^5	782.41	3.228	0.669	0.000441	12	0.13
61	2.3609×10^5	1.5652×10^5	782.25	3.100	0.672	0.000759	12	0.09
64	2.5333×10^5	1.3578×10^5	781.94	3.100	0.672	0.00215	12	0.11
67	2.8578×10^5	1.4215×10^5	781.25	2.876	0.667	0.00223	12	0.10
70 [†]	3.2658×10^5	436.34	811.16	2.532	0.696	0.0725	12	31.93
80	6.4448×10^5	54 658	773.12	2.690	0.669	0.00675	64	1.43
83	6.5083×10^5	57 106	773.07	2.690	0.669	0.000832	64	1.37
86	6.5993×10^5	56 532	772.92	2.690	0.670	5.22×10^{-6}	64	1.38
89	6.7172×10^5	57 304	772.76	2.689	0.670	3.35×10^{-6}	64	1.36
92	6.7909×10^5	58 481	772.10	2.691	0.671	5.57×10^{-6}	64	1.33
96	7.0438×10^5	54 257	771.77	2.687	0.671	3.62×10^{-6}	64	1.44
99	7.0822×10^5	55 039	771.39	2.687	0.672	9.5×10^{-7}	64	1.42
102	7.0963×10^5	64 176	771.09	2.818	0.672	0.00703	64	1.22
105	7.17×10^5	58 240	770.50	2.652	0.675	1.19×10^{-6}	64	1.34

S17.3 The excluded observation

The step-70 diagnostic reports $Q_{\text{FDTD}} = 436.3$ with a resonance-fit error of 0.0725—about two orders of magnitude above the surrounding scatter—together with a wavelength excursion to 811.16 nm, some 30 nm away from every neighbouring point. It is excluded from all aggregate statistics by the rule declared before the audit was run: resonance error above 0.01. That threshold excludes exactly this one point of 33; the next largest fit error in the record is 0.00703.

The exclusion is a fit-quality judgement, not a value judgement: it was applied without reference to the resulting Q , and the point is retained and marked in every figure rather than deleted. The most likely mechanism is a mode identification failure in the FDTD post-processing—the wavelength excursion and the fit error are consistent with the analysis latching onto a different resonance in the window—but we did not re-run it and do not claim to know.

S17.4 Protocol change

Diagnostics through step 67 used a 12 ps run time; those from step 80 onward used 64 ps. The change was made because the lengthening photon lifetime had overtaken the window: at step 61, $Q_{\text{FDTD}} = 156\,520$ implies $\tau \approx 130$ ps against a 12 ps window, i.e. $\mathcal{N} \approx 0.092$. The two segments are therefore different observables and every aggregate in the main text is reported for each separately as well as together.

S17.5 Geometry outcome

quantity	before		after	
	mean	range	mean	range
pitch (nm)	234.63	234.1–234.9	228.46	202.4–249.5
r_x (nm)	76.157	68.23–90.35	45.96	36.73–78.23
r_y (nm)	76.166	68.25–90.49	118.16	64.92–165.1
r_y/r_x	1.0001	0.9981–1.002	2.5613	1.36–3.112
$2r_y/W$	0.2738	0.2454–0.3253	0.4248	0.2334–0.5938
$1 - 2r_x/p$	0.3508	0.2301–0.4187	0.5986	0.373–0.6791
area (nm ²)	18 464	1.463e+04–2.568e+04	17 542	8743–4.058e+04

The mean longitudinal semi-axis contracts from 76.157 nm to 45.96 nm; the mean transverse semi-axis expands from 76.166 nm to 118.16 nm, with a maximum of 165.1 nm. The mean aspect ratio moves from 1.0001 to 2.5613, maximum 3.112. Near-circular holes become transverse slots spanning up to 0.5938 of the strip width, leaving dielectric rails of 113 nm per side against 187.6 nm initially. The mean hole *area* changes by a factor of only 0.9501: the air was reshaped, not added. The mean pitch contracts from 234.63 nm to 228.46 nm, and because pitches accumulate ((S10)) the outermost hole centre moves inward by 268 nm.

S17.6 Spectral and modal signature

The restructuring is accompanied by a blueshift of 21.8 nm, from 792.347 nm to 770.497 nm, and by a reduction of the effective mode volume from 5.047 to 2.652 $(\lambda/n)^3$, a factor of 0.525. The reported localization fraction is essentially unchanged, 0.679 to 0.675, so the mode did not migrate out of the cavity region.

Concentrating field energy into the air slots at fixed period would produce exactly this combination, and the geometry change is what one would design in order to obtain an air-type rather than a dielectric-type mode. We report it as a *characterized geometric and spectral restructuring consistent with* a change of mode character, and not as a demonstrated band-edge swap. Two reasons for that caution:

1. The available evidence is a correlation among hole shape, wavelength, and mode volume. The diagnostics that would settle the question—a band structure of the terminal mirror cell before and after, or a decomposition of the retained field energy between the dielectric and air regions—were not computed, and computing them would require re-analysis of raw field data that is outside the scope of this work.
2. A simple arithmetic argument points the other way. Evaluated on the mean pitch, $\lambda/2p$ is 1.689 at step 0 and 1.686 at step 105. The ratio of wavelength to period is essentially unchanged, so the blueshift is accounted for by the pitch contraction alone and cannot by itself be read as a jump to the other side of a gap.

S17.7 Independent evaluation and status

The exact retained endpoints:

	step 0	step 105
Q_{FDTD}	18 991.94	58 240.27
λ (nm)	792.347	770.497
Q_{ana}	8 958.74	717 003
$Q_{\text{ana}}/Q_{\text{FDTD}}$	0.472	12.31
$V_{\text{eff}} (\lambda/n)^3$	5.047	2.652
run time (ps)	12	64
$\mathcal{N}$	0.751	1.34

giving a gain of $3.0666\times$. The step-0 value is the endpoint of *this* trajectory; a Q of 20 942.06 from a separately qualified seed lineage exists in the record and must not be substituted for it, since doing so would change the quoted gain without any corresponding physical claim.

PROVISIONAL, for two independent reasons. The endpoints were not analyzed under matched settings—12 ps against 64 ps—and the predeclared panel of §S15 has not been executed. The mismatch runs in the direction that inflates the step-0 value and therefore deflates the quoted gain, but a mismatch in a favourable direction is still a mismatch.

S18 Rejected approximations and failures

A method paper that reports only what worked is not reproducible. This section catalogues the approaches that were tried and rejected, in the order they were encountered, with the symptom that exposed each one. Several produced plausible, converged-looking numbers, which is why they are worth recording.

S18.1 Few-mode and closed bases

Approach. Expand the cavity field in a handful of guided modes of the unpatterned beam, or in the modes of a domain terminated by an absorber, and read the linewidth from the resulting eigenvalue.

Symptom. A numerically precise Q that is essentially arbitrary. The few retained modes reproduce the near field well, so the resonance wavelength converges convincingly, while the tiny continuum coupling that actually determines $\text{Im } \tilde{\omega}$ is absent or replaced by an artefact of the absorber.

Resolution. The full guided-plus-continuum Green tensor of §S5. The imaginary part of the pole is generated by the continuum integral in (S35); no artificial linewidth is assigned to any mode anywhere in the construction.

S18.2 Unregularized collocation determinant

Approach. Sample the real-space dyadic Green kernel at quadrature nodes and take the determinant of $I - K$ directly.

Symptom. A determinant whose value depends on the coincident-point exclusion rule and which converges under refinement to something that is not the Fredholm determinant.

Resolution. Keep the longitudinal projector in momentum space and use a regularized determinant, §S4. The self term is then never evaluated as a singular real-space integral and no depolarization factor is guessed.

S18.3 Real-axis quadrature at complex frequency

Approach. Build the longitudinal quadrature on the real β axis, then substitute $\text{Im } \omega < 0$ to find the pole.

Symptom. A converged root, sometimes even passive, that is not the resonance: the substitution has silently continued onto a different sheet.

Resolution. The frozen scaled contour of (S45), moved along an explicit homotopy with singularity tracking, plus the closed-form orientation test (S46) as a precondition on any contour implementation.

S18.4 Hermitian conjugate in place of transpose

Approach. Use $C(\beta)^\dagger$ as the left factor in (S56), by analogy with a real-frequency energy inner product.

Symptom. Perfect agreement on the real-frequency boundary, then a continuation that does not converge to anything stable, and a loss of the complex symmetry that reciprocity requires.

Resolution. $C(-\beta)^T$, §S8. Complex symmetry of the assembled operator is now a running numerical check rather than an assumption.

S18.5 Layered cavity without feed-pole subtraction

Approach. Integrate (S56) for the multilayer device with the guided feed pole left in the integrand and the contour routed around it.

Symptom. A quality factor of order 80—three orders of magnitude below any plausible value—together with quadrature error that did not respond to node refinement at affordable cost.

Resolution. The Laurent extraction and exact real-space restoration of §S10. The rejected result is retained in the record and is superseded by the accepted pole $\tilde{\omega}/\omega_0 = 0.9992524494 + -6.871739 \times 10^{-5}i$, giving $Q_{\text{ana}} = 7270.74$ at 796.28805 nm against $Q_{\text{FDTD}} = 6176.55$ at 796.6645 nm: a Q disagreement of 17.72 % and a wavelength disagreement of -0.0473 %.

S18.6 The flat direction that looked like a $3.6\times$ gain

This is the most consequential failure in the record and the clearest single argument for the operator construction of this work.

Setting. A saved geometry direction from an early multilayer optimizer run, evaluated at iterations 0, 15, and 30.

Three verdicts.

evaluator	endpoint ratio
independent FDTD	1.0007628
pre-repair analytical operator	3.6045899
pole-subtracted analytical operator	1.0000101

The direction was flat. The pre-repair operator reported a $3.6\times$ improvement along it, which had it not been caught would have been an optimization campaign spent climbing an artefact. The repaired operator reproduces the flat verdict to five decimal places. Nothing was fitted or tuned to achieve this: the difference between the two analytical numbers is entirely the feed-pole treatment of §S10.

S18.7 Mesh-motion blockade

Approach. Apply the boundary velocity of (S69) to a meshed geometry with polygon repair, signed-distance reinitialization, and remeshing between steps.

Symptom. The line search stalled. The geometry actually delivered to the solver did not approach the previous geometry as the requested step went to zero, violating (S70), so no step size was small enough to produce a measurable, predictable change.

Resolution. The exact analytic parameterization of §S3, in which a step of size α moves the boundary by exactly α times the requested velocity. The mesh-based trajectory is retained as a counterexample.

S18.8 Short-ringdown quality factors

Approach. Fit Q from a 12 ps ringdown along a trajectory whose photon lifetime is growing.

Symptom. At step 61, $Q_{\text{FDTD}} = 156\,520$ with $\mathcal{N} \approx 0.092$; the subsequent 64 ps analyses of geometries with strictly higher analytical Q return values near 58 240.27.

Resolution. Partial. The protocol was lengthened mid-campaign, which is why the two segments must be treated separately, and the adequacy audit of §S15 now quantifies the problem for every retained point. The full resolution is the matched panel, with the run-time amendment recorded there.

S18.9 Objectives that have not passed their own gates

Effective mode volume and coupling derivatives were implemented but have not passed the algebraic and local-surrogate tests of §S12 to the standard the Q derivative has. They are therefore excluded from the quasi-Newton objective entirely and appear in this work only as reported diagnostics. Quality factor, mode volume, and coupling do not automatically share a trust level, and treating them as if they did is a failure mode in its own right.

S19 Reproducibility and provenance

S19.1 How the numbers reach this document

No numeric value in the main text or in this document is typed into the L^AT_EX source. Three generators read the retained evidence and emit machine records; two exporters turn those records into L^AT_EX macros and tables; the manuscripts reference the macros.

generator	output	used by
ringdown_adequacy.py	ringdown-adequacy.json	§S15; main-text ringdown table
geometry_metrics.py	geometry-metrics.json	§S16, §S17
surrogate_fidelity.py	surrogate-fidelity.json	main-text fidelity table
analytical_validation.py	analytical-validation.json	§S14
export_values.py	values.tex	every quoted number
export_tables.py	tables/*.tex	every long table

Two consequences are worth stating. First, when a matched panel later replaces a provisional endpoint, the JSON changes, the exporters re-run, and the prose is untouched—there is exactly one substitution point. Second, `analytical_validation.py` does not merely copy values out of the historical gate records: it declares, for each value, the literal string that must appear in the cited file, and it fails rather than emitting anything if that string is absent. The manuscript therefore cannot drift away from the record it cites.

S19.2 Statistical rules declared in advance

Two rules were fixed before the corresponding audits were run and are stated in the source of the generators that apply them.

Fit-quality exclusion.

An FDTD observation is disqualified when its reported resonance error exceeds 0.01. Applied to the retained trajectory this excludes exactly one point of 33 (§S17). The rule refers only to fit quality and never to the resulting quality factor.

Adequacy threshold.

A ringdown fit spanning fewer than one amplitude lifetime is flagged. This is a flag, not an exclusion: all such points are retained, marked, and reported.

S19.3 Provenance of the theory sections

Sections S3–S12 are written self-contained and in one notation. Each is nonetheless a rewrite of a specific retained derivation phase, and the mapping is below. The final column is checked at build time.

SI	topic	retained derivation, under <code>derivations/</code>	present
S3	exact parameterized dielectric	<code>quan_loncar/fryett_q_series_phase1</code>	✓
S4	Lippmann–Schwinger, $\det_{\text{reg}}$	<code>quan_loncar/fryett_q_series_phase1</code>	✓
S5	homogeneous reference Green tensor	<code>quan_loncar/fryett_q_series_phase2</code>	✓
S6	stratified reference Green tensor	<code>device23/multilayer_nanobeam_phase1</code>	✓
S7	Sommerfeld contour and quadrature	<code>quan_loncar/fryett_q_series_phase3</code>	✓
S8	Zernike–Legendre aperture basis	<code>quan_loncar/fryett_q_series_phase4</code>	✓
S9	symmetry reduction, array assembly	<code>quan_loncar/fryett_q_series_phase6</code>	✓
S10	feed-pole extraction and residues	<code>device23/multilayer_nanobeam_phase7</code>	✓
S11	continuation, pole search, mode identity	<code>quan_loncar/fryett_q_series_phase7</code>	✓
S12	implicit pole and Q derivatives	<code>qnm-boundary-gradient.md</code>	✓

Two notation hazards are worth repeating here, because a reader following the mapping into the raw record will meet them immediately. The air-clad lineage was developed against the encapsulated SiN benchmark of Fryett *et al.* [18]— $W = 450$ nm, $T = 330$ nm, $n_b = 2.0$, $n_c = 1.47$, $\lambda \approx 740$ nm—and not

against the production air-clad silicon device (§S16). And several multilayer phases describe a superseded 25-hole, 75-coordinate device rather than the production 42-hole, 126-coordinate one (§S17). The mathematics is device-independent; the constants are not.

S19.4 Raw artifacts

Field arrays from the FDTD campaigns total roughly 11 GB and are retained outside version control, indexed by absolute source path, byte size, and SHA-256 checksum. The compact records used in this work—endpoint summaries, the trajectory series, the frozen geometry vectors, and the diagnostic metadata—are all in version control and are what every generator above reads. Reproducing the analysis in this document therefore requires no access to the raw arrays; reproducing the FDTD numbers themselves does, or requires re-running the simulations.

S19.5 Build

A single command audits the claims ledger and the provenance index, runs the test suite, regenerates every figure and table, compiles both documents with a pinned $\text{\TeX}$ distribution, and builds the documentation site. The tests pin the exact values quoted here, so a change to any record that is not accompanied by a change to the corresponding assertion fails the build. No build command is able to submit a paid simulation; the revalidation plan is stored with an explicit flag recording that submission is not permitted by the build, and the panel status is predeclared not submitted.

References

- [1] Philippe Lalanne, Wei Yan, Kevin Vynck, Christophe Sauvan, and Jean-Paul Hugonin. Light interaction with photonic and plasmonic resonances. *Laser & Photonics Reviews*, 12(5):1700113, 2018.
- [2] Philip Trost Kristensen, Kathrin Herrmann, Francesco Intravaia, and Kurt Busch. Modeling electromagnetic resonators using quasinormal modes. *Advances in Optics and Photonics*, 12(3):612–708, 2020.
- [3] P. T. Leung, S. Y. Liu, and K. Young. Completeness and orthogonality of quasinormal modes in leaky optical cavities. *Physical Review A*, 49:3057–3067, 1994.
- [4] Stefan Güttel and Françoise Tisseur. The nonlinear eigenvalue problem. *Acta Numerica*, 26:1–94, 2017.
- [5] Alexander S. Markus. Introduction to the spectral theory of polynomial operator pencils. *Translations of Mathematical Monographs*, 71, 1988.
- [6] Arthur D. Yaghjian. Electric dyadic Green’s functions in the source region. *Proceedings of the IEEE*, 68(2):248–263, 1980.
- [7] Barry Simon. *Trace Ideals and Their Applications*. American Mathematical Society, 2nd edition, 2005.
- [8] Frank W. J. Olver, Daniel W. Lozier, Ronald F. Boisvert, and Charles W. Clark. *NIST Handbook of Mathematical Functions*. Cambridge University Press, 2010.
- [9] Dennis P. Nyquist, D. R. Infante, and J. S. Bagby. Orthogonality and amplitude spectrum of radiation modes along open-boundary waveguides. *Journal of the Optical Society of America*, 71(1):49–54, 1981.
- [10] Max Born and Emil Wolf. *Principles of optics*. 1999.
- [11] Lloyd N. Trefethen and J. A. C. Weideman. The exponentially convergent trapezoidal rule. *SIAM Review*, 56(3):385–458, 2014.

- [12] Steven G. Johnson, M. Ibanescu, M. A. Skorobogatiy, O. Weisberg, J. D. Joannopoulos, and Y. Fink. Perturbation theory for Maxwell’s equations with shifting material boundaries. *Physical Review E*, 65:066611, 2002.
- [13] Jan Sokółowski and Jean-Paul Zolésio. *Introduction to Shape Optimization: Shape Sensitivity Analysis*. Springer, 1992.
- [14] Allen Taflove and Susan C. Hagness. *Computational Electrodynamics: The Finite-Difference Time-Domain Method*. Artech House, 3rd edition, 2005.
- [15] Jean-Pierre Bérenger. A perfectly matched layer for the absorption of electromagnetic waves. *Journal of Computational Physics*, 114(2):185–200, 1994.
- [16] Flexcompute. Tidy3D: GPU-accelerated finite-difference time-domain solver. <https://www.flexcompute.com/tidy3d/>, 2024.
- [17] Qimin Quan and Marko Lončar. Deterministic design of wavelength scale, ultra-high Q photonic crystal nanobeam cavities. *Optics Express*, 19(19):18529–18542, 2011.
- [18] Taylor K. Fryett, Yueyang Chen, James Whitehead, Zane M. Peycke, Xiaodong Xu, and Arka Majumdar. Encapsulated silicon nitride nanobeam cavity for hybrid nanophotonics. *ACS Photonics*, 5(6):2176–2181, 2018.