Architecture-Dependent Decoherence Suppression in Passive Quantum Networks: Irreducible Channel Mixing, Squared Rate Gaps, and the Price in Dwell Time

Lluis Eriksson

10 August 2026

Abstract

Can passive reservoir filtering suppress decoherence more strongly when its channel mixing is irreducible, even after passband responses and rational complexity are fixed? We give an explicit affirmative separation, together with the resource that pays for it. For every integer $S \geq 5$ we construct a causal three-signal/three-loss-port inner transfer whose signal block G_S is a rational Schur function with numerator degree at most $S+4$ and common denominator degree S. It obeys $3(S-1)$ exact delayed tangential calibrations in a pass arc; every three calibration signatures span $\mathbb{C}^3$. On two stop arcs,

$$\sup \|G_S\|_{\mathrm{op}} \leq \varepsilon_S := \left(\frac{\pi S}{6} + \frac{1}{2}\right) \mathrm{e}^{-9S/20}.$$

In contrast, every transfer of the same bidegree that has a constant nontrivial reducing channel and obeys the same calibrations has stopband norm at least one. The proof is a scalar root-count obstruction exposed by the full-spark signatures. We strengthen this discontinuous-looking statement: for calibration defect δ and sampled reducing-line defect β, the comparator leakage is at least $[1 - C_S(\delta + \beta)]_+$, where C_S is an explicit finite-matrix constant. A scalar Schur construction then proves that every valid C_S must grow at least as $2\mathrm{e}^{9S/20}(1 + o(1))$; uniform robustness is impossible for these calibrations. The lossless completion uses only three scalar Blaschke all-pass elements, balanced interferometers and paraunitary delays, so physical passivity is verified directly rather than inferred from an unconstrained interpolation. For a bath spectral density $j_- I \preceq J \preceq j_+ I$, the filtered Kossakowski matrix therefore has norm at most $j_+ \varepsilon_S^2$, whereas every reducing comparator retains a directional rate at least j_-. In a closed Markov pure-dephasing model the three auxiliary vacuum ports contribute exactly the architecture-independent baseline $\kappa_0 I$; hence the observable total-rate advantage is at least $(\kappa_0 + j_-)/(\kappa_0 + j_+ \varepsilon_S^2)$ and saturates at the vacuum floor. The separation is not free: if a finite Blaschke product of degree S yields signal leakage at most $\varepsilon < 1$ on an arc of length L, its peak group delay is at least $LS/(16 \arcsin \varepsilon)$. Our family has exponentially large delay. Thus under a peak-delay budget D, the rate-improvement factor is at most $\csc^2(LS/(16D))$, asymptotically quadratic in D/S. Irreducible passive mixing can therefore move a coherence-maintenance boundary, but it moves the required resource into long-lived passive memory rather than erasing it. Exact arithmetic, dense-grid unitarity audits and all figures are reproduced by the linked public source package.

1 Question, result, and non-claims

Quantum filter functions translate control protocols into frequency-resolved noise response [1, 2], and passive linear quantum systems admit network and realization theories that preserve the canonical commutation relations [3, 4, 5]. Tangential interpolation can itself be made physical-realizability- and passivity-preserving [6]. Blaschke–Potapov factors have also been used to identify trapped modes in delayed passive quantum networks [7]. None of those ingredients is claimed as new.

The narrower question here is architectural. Suppose that a multichannel filter must match many directional responses exactly, has a fixed rational bidegree, and remains passive. Is a device whose channels split after one frequency-independent rotation necessarily as capable as one whose channel geometry changes with frequency? The answer is no. We prove an exponential signal-level gap between an explicit irreducible signal block and every constant-channel reducing comparator. Passing that gap through a uniformly nondegenerate bath spectral density squares the suppression exponent. We then prove a delay lower bound that prevents the architecture theorem from being misread as a free thermodynamic advantage.

This result develops the resource-boundary viewpoint of [8] in a direction orthogonal to spatial buffer thickness. The protected object is a frequency-resolved coupling channel. The filter can move the rate boundary by changing its channel architecture, but poles near the unit circle store the resource as dwell time. The paper makes no claim to derive the Born rule, select outcomes, solve an interacting many-body buffer, or show that a small thermodynamic lower bound is achievable.

The contributions are:

1. an explicit six-port causal inner completion, with no appeal to a black-box synthesis theorem;
2. a fixed-bidegree irreducible/reducible separation under exact full-spark tangential calibrations;
3. an approximate-calibration theorem with a computable conditioning constant, and a matching exponential obstruction to uniform robustness;
4. a squared separation for filtered Kossakowski matrices and a closed Ramsey model in which all vacuum ports are accounted for;
5. a general delay-price theorem for the scalar all-pass mechanism;
6. an exact no-go showing why the earlier real-symmetric quadratic certificate cannot be made globally passive by scalar rescaling; and
7. a reproducible ledger linking every numerical statement to code [9].

2 Passive transfer model and comparison class

We use the unit-disc convention. A matrix transfer $\mathsf{S}(z)$ is *inner* when it is analytic for $|z| < 1$ and unitary almost everywhere on $\mathbb{T} = \{z : |z| = 1\}$. A principal block $G(z)$ of an inner transfer is Schur: $\|G(z)\|_{\text{op}} \leq 1$ in the disc and on its regular boundary points. In a passive quantum-optical interpretation, the omitted ports carry vacuum loss; unitarity of the full annihilation-operator transfer preserves the canonical commutation relations. Delays and stable all-pass factors are included, so the model is a discrete-time or sampled traveling-field network, not an instantaneous scattering matrix.

Let

$$\mathcal{I}_{\text{p}} = [-\pi/6, \pi/6], \qquad \mathcal{I}_{\text{s}} = [-\pi, -5\pi/6] \cup [5\pi/6, \pi].$$

Given boundary nodes $z_j = \mathrm{e}^{\mathrm{i}\theta_j}$ and unit signatures $v_j \in \mathbb{C}^3$, the prescribed response is the two-sample delay

$$Q(z_j)v_j = z_j^2 v_j. \tag{1}$$

This common delay makes the calibration compatible with causality and with a lossless pass response.

Definition 2.1 (Rational bidegree class)**.** For $S \geq 1$, let $\mathcal{R}_S$ be the set of 3×3 rational transfers $Q = A/d$, regular on $\mathbb{T}$, whose scalar common denominator satisfies $\deg d \leq S$ and whose numerator entries satisfy $\deg A_{ab} \leq S+4$. We take d to have no zeros on $\mathbb{T}$, after removing common cancellations. The class allows cancellations; using upper bounds rather than exact degrees only strengthens the comparator.

Definition 2.2 (Constant-channel reducibility)**.** A transfer Q is *reducing* if a nonzero proper subspace $E \subset \mathbb{C}^3$, independent of z, reduces $Q(z)$ at every regular point: both E and $E^\perp$ are invariant. In dimension three, either E or its orthogonal complement contains a reducing line $\mathbb{C}u$. Thus for a unit u there is a scalar rational branch q with

$$Q(z)u = q(z)u, \qquad Q(z)^*u = \overline{q(z)}u \quad (z \in \mathbb{T}).$$

This comparator is strictly larger than simultaneously diagonalizable or commuting channels: the complementary 2×2 block may be fully noncommutative and frequency dependent. The comparison therefore isolates a constant channel split, not mere pairwise commutativity.

Proposition 2.3 (Static-split characterization)**.** *A transfer Q has a constant reducing subspace if and only if there is one frequency-independent unitary R for which*

$$R^*Q(z)R = \begin{pmatrix} Q_1(z) & 0 \\ 0 & Q_2(z) \end{pmatrix}$$

at every regular point, with both blocks nonempty. The property is invariant under constant unitary changes of channel coordinates.

Proof. Choose orthonormal bases of the reducing subspace and its orthogonal complement and concatenate them into R. Invariance of both subspaces kills the two off-diagonal blocks. Conversely, the span of either coordinate block reduces the displayed transfer. Conjugating by another constant unitary just transports that subspace. □

Thus the comparator is precisely a network that separates into parallel channel groups after a static reciprocal change of coordinates; it is stable under static relabeling and parallel composition. A frequency-dependent diagonalizer is itself a dynamical filter with poles, delay and degree, so granting it for free would change the resource accounting. Distinct constant rotations on input and output define a broader two-sided factorization and are not covered: the present question concerns a persistent physical channel direction shared by injection, propagation and readout.

3 Why the earlier zero-phase witness cannot simply be made passive

The earlier matrix-polynomial construction [10] used four real-symmetric quadratic tangential constraints. It is tempting to rescale that witness until its norm is at most one. The following exact calculation shows that the idea fails before any quantum interpretation is attempted.

Let

$$x = (1/2, 1, 3/2, 7/2), \qquad v = ((1,2), (0,1), (-1,2), (-1,1)),$$

and impose $P(x_j)v_j = v_j$ on a real-symmetric quadratic P. Define

$$R(x) = R_0 + xR_1 + x^2R_2$$

with

$$R_0 = \begin{pmatrix} 5/6 & -1/3 \\ -1/3 & 5/6 \end{pmatrix}, \quad R_1 = \begin{pmatrix} -2/3 & -1/3 \\ -1/3 & -11/6 \end{pmatrix}, \quad R_2 = \begin{pmatrix} 2/3 & 2/3 \\ 2/3 & 1 \end{pmatrix}.$$

Proposition 3.1 (Exact Hermitian passivity no-go)**.** *Every feasible real-symmetric quadratic is $P_\tau = I + \tau R$. If $\|P_\tau(x)\|_{\mathrm{op}} \le 1$ for every real x, then $\tau = 0$.*

Proof. In the symmetric-entry coordinates $(a_k, b_k, c_k)_{k=0}^2$, the eight scalar calibration equations have rational rank eight. Direct substitution shows that $R(x_j)v_j = 0$, so the affine solution space is exactly $I + \tau R$. Now $R(0)$ is positive definite because its leading entry is 5/6 and its determinant is 7/12. Hence $\tau > 0$ gives an eigenvalue of $P_\tau(0)$ strictly above one. At $x = 5/2$,

$$R(5/2) = \begin{pmatrix} 10/3 & 3 \\ 3 & 5/2 \end{pmatrix}, \qquad \det R(5/2) = -2/3,$$

so $R(5/2)$ has a negative eigenvalue; $\tau < 0$ again makes an eigenvalue of $P_\tau(5/2)$ exceed one. Only $\tau = 0$ remains. □

The obstruction motivates the boundary-phase construction below. Its calibration is delayed rather than zero phase, its coefficients are complex causal transfer coefficients rather than Hermitian polynomials on the real line, and passivity is built into a lossless completion.

4 Explicit Blaschke–interferometer construction

Fix $\alpha = 9/20$ and $S \geq 5$. Put

$$r_S = 1 - e^{-\alpha S}, \qquad b_r(z) = \frac{z - r}{1 - rz}, \qquad B_S(z) = \eta_S b_{r_S}(z)^S, \tag{2}$$

where $\eta_S = 1$ for odd S and $\eta_S = -1$ for even S. Thus $B_S(-1) = -1$. Set $\delta_S = 1 - r_S$ and

$$\phi_g = (g-1)\delta_S, \qquad p_g(z) = \frac{1 + e^{i\phi_g} B_S(z)}{2}, \quad g = 0, 1, 2. \tag{3}$$

Let $F_{ab} = 3^{-1/2}\omega^{ab}$ with $\omega = e^{2\pi i/3}$, and define paraunitary mixers

$$U(z) = F \operatorname{diag}(1, z, z^2)F^*, \qquad U_{\text{rev}}(z) = F \operatorname{diag}(z^2, z, 1)F^*. \tag{4}$$

On $\mathbb{T}$, $U_{\text{rev}}(z) = z^2 U(z)^*$. The signal transfer is

$$G_S(z) = U(z) \operatorname{diag}(p_0(z), p_1(z), p_2(z)) U_{\text{rev}}(z). \tag{5}$$

4.1 Direct six-port lossless completion

For $B_g(z) = e^{i\phi_g} B_S(z)$ define the balanced Mach–Zehnder block

$$W_g(z) = \frac{1}{2}\begin{pmatrix} 1 + B_g(z) & 1 - B_g(z) \\ 1 - B_g(z) & 1 + B_g(z) \end{pmatrix} = H \operatorname{diag}(1, B_g(z))H, \quad H = 2^{-1/2}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}. \tag{6}$$

Each W_g is inner and has signal entry p_g. Let W place the three W_g blocks on the signal/loss pairs $(g, g+3)$. Then

$$\mathsf{S}_S(z) = \operatorname{diag}(U(z), I_3)\, W(z)\, \operatorname{diag}(U_{\text{rev}}(z), I_3) \tag{7}$$

is inner, and its upper-left 3×3 block is G_S.

Proposition 4.1 (Physical passivity by construction)**.** *S_S is a causal rational inner six-port transfer and G_S is a rational Schur transfer in $\mathcal{R}_S$.*

Proof. b_r is analytic in the unit disc and unimodular on $\mathbb{T}$. Hence every B_g, W_g, W, U, and U_{rev} is inner; products and block diagonals preserve innerness. The signal-block contraction follows from the unitarity of the boundary completion. The common denominator is $(1 - r_S z)^S$. The scalar p_g numerators have degree at most S, while the two mixers add at most four, so the numerator entries of G_S have degree at most $S + 4$. □

The factorization is also a component prescription: three repeated scalar all-pass elements, three balanced interferometers, a three-mode Fourier network, and integer delays. General passive-network realization results [5, 11] provide a broader context, but are not needed for the boundary unitarity identity used here.

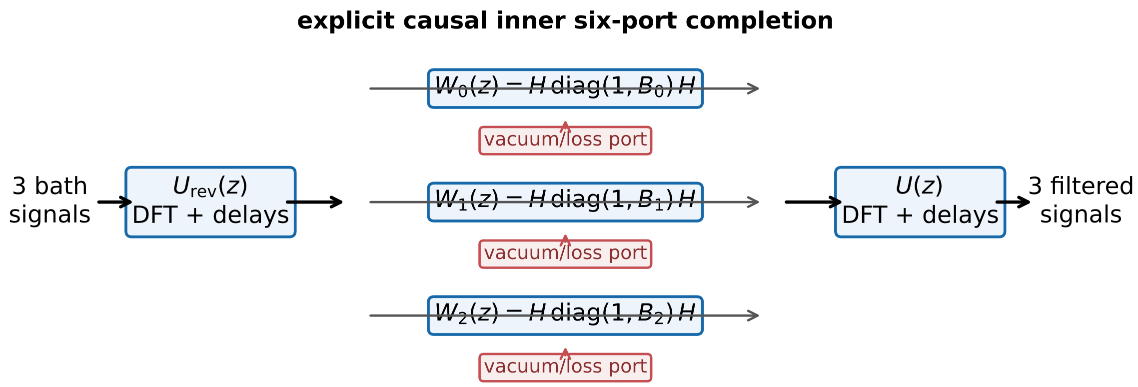

Figure 1: Component-level realization. The signal passes through a reverse paraunitary mixer, three parallel lossless Mach–Zehnder/all-pass blocks and a forward mixer. The three auxiliary inputs are vacuum/loss ports. Every displayed block is inner, so their cascade is inner.

4.2 Pass nodes and full-spark signatures

At a point where $B_g(z) = 1$, let

$$v_{g,z} = U(z)e_g. \tag{8}$$

Then $p_g(z) = 1$ and the identity $U_{\mathrm{rev}} = z^2 U^*$ gives

$$G_S(z)v_{g,z} = z^2 v_{g,z}. \tag{9}$$

Lemma 4.2 (Node count)**.** *For each $g \in \{0, 1, 2\}$, the open pass arc contains at least $S - 1$ distinct solutions of $B_g(z) = 1$.*

Proof. On the unit circle write $b_r(\mathrm{e}^{\mathrm{i}\theta}) = \mathrm{e}^{\mathrm{i}\beta_r(\theta)}$ with a continuous unwrapped phase on $(-\pi, \pi)$ and

$$\beta_r'(\theta) = \frac{1 - r^2}{1 - 2r\cos\theta + r^2} > 0.$$

For $r = r_S$, put $\beta_0 = \beta_r(\pi/6)$. The identity

$$\beta_0 = 2\arctan\Big((2\mathrm{e}^{9S/20} - 1)\tan\frac{\pi}{12}\Big)$$

and $\cot x < 1/x$ show that $\beta_0 > \pi(1 - 1/S)$ for every $S \geq 5$: the sufficient inequality $(2\mathrm{e}^{9S/20} - 1)\tan(\pi/12) > 2S/\pi$ holds at $S = 5$ and its left-minus-right derivative is positive thereafter. More directly, the inverse formula

$$\theta = 2\arctan\left(\frac{1 - r}{1 + r}\tan\frac{\beta}{2}\right)$$

maps the interval $(-\beta_0, \beta_0)$ into the pass arc. The allowed interval for the integer winding number in $S\beta + \arg\eta_S + \phi_g = 0 \pmod{2\pi}$ has length $S\beta_0/\pi > S - 1$. Every open real interval of length greater than $S - 1$ contains at least $S - 1$ integers, independently of the phase offset. Hence at least $S - 1$ roots are interior. □

Select any $S - 1$ nodes from each group and call the resulting set $\mathcal{Z}_S$; its size is $M = 3(S - 1)$. These signatures have a stronger property than spanning collectively.

Lemma 4.3 (Full spark)**.** *Every three distinct signatures in $\{v_{g,z} : (g, z) \in \mathcal{Z}_S\}$ are linearly independent.*

Proof. Multiplication by F^* does not change independence, and

$$F^* v_{g,z} = 3^{-1/2}(1, t, t^2)^\top, \qquad t = \omega^{-g} z.$$

The three pass-node groups are rotated by $0, -2\pi/3, -4\pi/3$; their arcs are disjoint, so distinct signatures give distinct t. The determinant of any three columns is the nonzero Vandermonde product $3^{-3/2}(t_2 - t_1)(t_3 - t_1)(t_3 - t_2)$. $\square$

4.3 Stopband leakage

Lemma 4.4 (Analytic stop envelope)**.** *On $\mathcal{I}_\mathrm{s}$,*

$$\|G_S(\mathrm{e}^{\mathrm{i}\theta})\|_{\mathrm{op}} \le \varepsilon_S := \left(\frac{\pi S}{6} + \frac{1}{2}\right) \mathrm{e}^{-9S/20}. \tag{10}$$

Proof. The mixers are unitary on $\mathbb{T}$, so the norm is $\max_g |p_g|$. The all-pass phase derivative is S times the Poisson kernel. On either stop arc, $\cos\theta \le -\sqrt{3}/2$, whence

$$\frac{1 - r^2}{1 - 2r\cos\theta + r^2} \le 1 - r^2.$$

Starting from $B_S(-1) = -1$, its phase deviates from π by at most $S(1 - r^2)\pi/6$. Including $|\phi_g| \le 1 - r$ and using $|(1 + \mathrm{e}^{\mathrm{i}(\pi + x)})/2| \le |x|/2$ gives

$$|p_g| \le \frac{\pi S}{12}(1 - r^2) + \frac{1}{2}(1 - r) \le \left(\frac{\pi S}{6} + \frac{1}{2}\right)(1 - r).$$

Substitute $1 - r = \mathrm{e}^{-\alpha S}$. $\square$

5 The architecture-separation theorem

We can now state the central result.

Theorem 5.1 (Passive irreducible/reducing separation)**.** *Fix $S \ge 5$ and $\alpha = 9/20$. There are $M = 3(S - 1)$ pass nodes in $\mathcal{I}_\mathrm{p}$ and full-spark signatures such that:*

1. *the signal block $G_S \in \mathcal{R}_S$ has the explicit six-port inner completion* (7)*, obeys all calibrations* (1)*, and has stop leakage at most ε_S;*
2. *every reducing $Q \in \mathcal{R}_S$ obeying the same calibrations satisfies*

$$\|Q(z)\|_{\mathrm{op}} \ge 1 \qquad (z \in \mathcal{I}_\mathrm{s}); \tag{11}$$

3. *if such a comparator is Schur, equality holds in* (11) *throughout the stopband;*
4. *whenever $\varepsilon_S < 1$, G_S is irreducible.*

Thus the signal-level architecture ratio is at least $1/\varepsilon_S$ and is exponential up to the displayed polynomial prefactor.

Proof. The constructive statements follow from Theorems 4.1 to 4.4. It remains to prove the comparator obstruction. Let $\mathbb{C}u$ be a reducing line of Q and let q be its scalar branch. Left-projecting each calibration with u^* gives

$$(q(z_j) - z_j^2)\langle u, v_j \rangle = 0. \tag{12}$$

Full spark implies that at most two signatures lie in the plane $u^\perp$; otherwise three would be dependent. Consequently, $q - z^2$ vanishes at at least

$$M - 2 = 3S - 5$$

distinct boundary points. After placing q over the common denominator d, the numerator of $q - z^2$ has degree at most $S + 4$. Since $3S - 5 > S + 4$ for $S \geq 5$, the numerator is identically zero. Therefore $q(z) = z^2$ wherever regular, and $\|Q(z)\|_{\mathrm{op}} \geq \|Q(z)u\| = 1$. A Schur comparator supplies the reverse inequality. Finally, if G_S were reducing, the same result would force its stop norm to be at least one, contradicting $\varepsilon_S < 1$. □

Remark 5.2 (What is held fixed). Both sides have three signal channels, identical delayed directional calibrations and the same rational bidegree ceiling. The irreducible side is also required to possess an explicit passive completion. The lower bound is proved for a larger comparison class that need not be passive; imposing passivity only turns its lower bound into equality.

5.1 Approximate calibration and approximate reduction

The exact obstruction has a quantitative counterpart. Put $N = S + 4$ and $K = M - 2 = 3S - 5$, and define the full-spark singular margin

$$\gamma_S = \min_{|J|=3} \sigma_{\min}[v_j]_{j \in J} > 0. \tag{13}$$

For an erasure set $E \subset \{1, \ldots, M\}$ of size two, let $V_E = (z_j^k)_{j \notin E,\, 0 \leq k \leq N}$ and set

$$\chi_S = \min_{|E|=2} \sigma_{\min}(V_E) > 0. \tag{14}$$

Every V_E has full column rank because it contains at least $N + 1$ distinct Vandermonde rows. Finally take $R = 8(S + 1)$ fixed midpoint samples ξ_ℓ split equally across the two stop arcs, form $W_S = (\xi_\ell^k)_{1 \leq \ell \leq R,\, 0 \leq k \leq S}$, and put

$$\nu_S = \frac{\sigma_{\min}(W_S)}{\sqrt{R(S+1)}} > 0. \tag{15}$$

For $Q \in \mathcal{R}_S$ and a unit vector u, define the calibration and sampled left-reduction defects

$$\Delta(Q) = \max_j \|Q(z_j)v_j - z_j^2 v_j\|, \quad \beta_u(Q) = \max_j \|(I - uu^*)Q(z_j)^* u\|. \tag{16}$$

The second quantity vanishes for the reducing line used in Theorem 5.1; allowing it to be small also covers approximate one-way decoupling at the measured nodes.

Theorem 5.3 (Quantitative approximate obstruction)**.** *For every $Q \in \mathcal{R}_S$ and every unit u,*

$$\sup_{\theta \in \mathcal{I}_{\mathrm{s}}} \|Q(\mathrm{e}^{\mathrm{i}\theta})\|_{\mathrm{op}} \geq \big[1 - C_S\big(\Delta(Q) + \beta_u(Q)\big)\big]_+, \qquad C_S = \frac{\sqrt{3K(N+1)}}{\gamma_S \chi_S \nu_S}, \tag{17}$$

where $[x]_+ = \max\{x, 0\}$. In particular, this is a finite-error theorem for every exactly reducing comparator by taking $\beta_u(Q) = 0$.

Proof. Let $q = u^* Q u$. Order the overlaps $|\langle u, v_j \rangle|$ and erase the two smallest. Applied to the triple of three smallest overlaps, (13) implies that the third, and hence every retained overlap, is at least $\gamma_S/\sqrt{3}$. Writing $Q^* u = \overline{q}\, u + w$ at a calibration node gives

$$|(q(z_j) - z_j^2)\langle u, v_j \rangle| \leq \Delta(Q) + \beta_u(Q).$$

Normalize the common denominator so that $\sup_{\mathbb{T}} |d| = 1$, write $q = a/d$ and $h = a - z^2 d$. Then $\deg h \leq N$ and, on the K retained nodes,

$$|h(z_j)| \leq \sqrt{3}\,[\Delta(Q) + \beta_u(Q)]/\gamma_S.$$

The least-squares bound from (14) gives $\|c_h\|_2 \le \sqrt{3K}[\Delta(Q)+\beta_u(Q)]/(\gamma_S\chi_S)$, so $\sup_{\mathbb{T}}|h| \le \sqrt{N+1}\|c_h\|_2$.

It remains to prevent the denominator from hiding this estimate. If c_d is its padded coefficient vector, then $1 = \sup_{\mathbb{T}}|d| \le \sqrt{S+1}\|c_d\|_2$. Hence (15) guarantees a fixed stop sample ξ_ℓ with $|d(\xi_\ell)| \ge \nu_S$. At that point, $|q-\xi_\ell^2| = |h/d| \le C_S[\Delta(Q)+\beta_u(Q)]$, and therefore $\|Q(\xi_\ell)\|_{\mathrm{op}} \ge |u^*Q(\xi_\ell)u| \ge 1 - C_S[\Delta(Q)+\beta_u(Q)]$. Taking the positive part proves the claim. □

The price of a denominator-free statement is visible in C_S. More importantly, the deterioration cannot be removed by a sharper proof.

Theorem 5.4 (No uniform robustness for these calibrations)**.** *Let C_S' be any constant for which $\sup_{\mathcal{I}_s}\|Q\|_{\mathrm{op}} \ge [1-C_S'\Delta(Q)]_+$ holds for all reducing Schur transfers in $\mathcal{R}_S$. Then*

$$C_S' \ge \frac{[1-\varepsilon_S]_+}{\sin(\mathrm{e}^{-9S/20}/2)} = 2\mathrm{e}^{9S/20}(1+o(1)). \tag{18}$$

Proof. Take the scalar transfer $Q_{\mathrm{c}}(z) = z^2(1+B_S(z))I_3/2$. It is Schur, reducing and belongs to $\mathcal{R}_S$. At a calibration from group g, one has $B_S(z_j) = \mathrm{e}^{-\mathrm{i}\phi_g}$, so its maximum calibration defect is exactly $\sin(\mathrm{e}^{-9S/20}/2)$. The central branch is one of the three scalar arms in the construction, and Theorem 4.4 gives $\sup_{\mathcal{I}_s}\|Q_{\mathrm{c}}\| \le \varepsilon_S$. Substitution in the assumed inequality yields (18); the asymptotic follows from $\varepsilon_S \to 0$ and $\sin x \sim x$. □

Together, Theorems 5.3 and 5.4 replace a qualitative fragility caveat by a precise statement: finite errors are controlled at every fixed S, but the present increasingly clustered calibration family cannot support an S-uniform tolerance theorem.

6 Transfer to bath spectra and decoherence rates

Let $J(\omega)$ be the 3×3 bath spectral-density matrix in a protected frequency band. Assume the uniform nondegeneracy condition

$$j_- I \preceq J(\omega) \preceq j_+ I, \qquad 0 < j_- \le j_+ < \infty. \tag{19}$$

For a linear passive prefilter T, the output noise spectrum is

$$\Gamma_T(\omega) = T(\mathrm{e}^{\mathrm{i}\theta(\omega)})J(\omega)T(\mathrm{e}^{\mathrm{i}\theta(\omega)})^*. \tag{20}$$

In a weak-coupling Davies construction, Γ_T is the Kossakowski matrix sampled at the relevant Bohr frequency [12]; in stationary Gaussian pure dephasing it is directly the quadratic rate matrix. The following statement is linear algebra and therefore does not depend on a particular master-equation normalization.

Corollary 6.1 (Squared Kossakowski gap)**.** *On the stopband, the constructed signal block obeys*

$$\|\Gamma_{G_S}(\omega)\|_{\mathrm{op}} \le j_+\varepsilon_S^2. \tag{21}$$

For every reducing calibrated comparator $Q \in \mathcal{R}_S$, there is a fixed unit direction u such that

$$u^*\Gamma_Q(\omega)u \ge j_- \tag{22}$$

throughout the stopband. Hence the worst retained directional rate divided by the constructed worst-case rate is at least $j_-/(j_+\varepsilon_S^2)$.

Proof. Equation (21) follows from $\|TJT^*\| \le \|T\|^2\|J\|$. For the comparator, the proof of Theorem 5.1 gives $Q^*u = \bar{z}^2 u$ on $\mathbb{T}$. Thus

$$u^*QJQ^*u = (Q^*u)^*J(Q^*u) \ge j_-\|Q^*u\|^2 = j_-.$$

□

For a transition whose three coupling matrix elements form a vector c, its Davies or pure-dephasing quadratic susceptibility is proportional to $c^*\Gamma_T c$. Consequently Theorem 6.1 gives an exact rate separation for any model in which the relevant transition can probe the retained direction u. It also gives an architecture-independent upper bound for all unit coupling vectors under G_S.

In pointer-basis pure dephasing, the coherence lifetime is inversely proportional to the corresponding quadratic rate. The construction therefore permits a lifetime gain of order ε_S^{-2} relative to the retained comparator direction, subject to the Markov, stationarity and spectral-bounds assumptions above. Applying a maintenance inequality of the form $P_{\mathrm{extra}} \geq k_B T \dot{C}_{\mathrm{loss}}$, as carefully delimited in [8], transfers the rate comparison to the *input of the lower bound.* A smaller lower bound does not prove an achievable low-power maintenance protocol, and we do not use it that way.

6.1 A closed Ramsey model with the vacuum ports included

The signal-only congruence (20) does not by itself specify the fluctuations entering the auxiliary ports. Here is a minimal model in which they are included rather than appended as an unknown correction. Write the signal output row of the full six-port inner transfer as $[T\ L]$ and take the six-input symmetrized white-noise covariance to be

$$V_{\mathrm{in}} = \frac{1}{2} I_6 + \begin{pmatrix} J & 0 \\ 0 & 0 \end{pmatrix}. \tag{23}$$

Thus every port contains vacuum, while the first three carry the additional stationary Gaussian bath spectrum J. Couple a pointer contrast $c \in \mathbb{C}^3$ linearly to the three output quadratures in the Markov pure-dephasing limit, and absorb the convention-dependent coupling prefactor into the rate units.

Proposition 6.2 (Vacuum-complete Ramsey rate)**.** *For every passive six-port completion,*

$$V_{\mathrm{out,sig}} = \frac{1}{2} I_3 + TJT^*, \qquad \kappa_T(c) = \kappa_0 \|c\|^2 + c^* TJT^* c, \tag{24}$$

where κ_0 is the common vacuum Ramsey rate in the chosen normalization. Consequently the constructed network and every reducing calibrated comparator $Q \in \mathcal{R}_S$ that admits such a passive completion obey

$$\sup_{\|c\|=1} \kappa_{G_S}(c) \leq \kappa_0 + j_+ \varepsilon_S^2, \qquad \kappa_Q(u) \geq \kappa_0 + j_-, \tag{25}$$

and their certified total-rate advantage is at least

$$\frac{\kappa_0 + j_-}{\kappa_0 + j_+ \varepsilon_S^2}. \tag{26}$$

Proof. Unitarity of the signal output row gives $TT^* + LL^* = I_3$. Applying $[T\ L]$ on both sides of (23) therefore yields

$$\tfrac{1}{2}(TT^* + LL^*) + TJT^* = \tfrac{1}{2} I_3 + TJT^*.$$

Markov Gaussian pure dephasing turns the quadratic covariance in direction c into the Ramsey exponent, which gives (24). The two rate bounds are then Theorem 6.1 plus the same κ_0 term. $\square$

This model makes two experimentally distinct claims. Switching off the excess bath preparation ($J = 0$) measures κ_0 directly; switching it on measures the excess rate $\kappa - \kappa_0$. The latter retains the squared exponential architecture gap, while the total lifetime benefit saturates at $1 + j_-/\kappa_0$ as $S \to \infty$ when $j_- = j_+$. Vacuum therefore limits the observable gain but does not invalidate or ambiguously renormalize it. The claim remains deliberately within a linear, Gaussian and Markov setting.

7 The hidden resource: a delay-price theorem

Exponential suppression in Theorem 5.1 is obtained by taking r_S exponentially close to one. This creates long-lived internal modes. The next result shows that this is not an accident of our parameterization.

Let

$$B(z) = \xi \prod_{k=1}^{S} \frac{z - a_k}{1 - \bar{a}_k z}, \qquad |\xi| = 1, \quad |a_k| < 1,$$

and $p = (1 + B)/2$. On the unit circle let $\tau_B(\theta)$ denote the unwrapped phase derivative, or group delay in sample units:

$$\tau_B(\theta) = \sum_{k=1}^{S} \frac{1 - |a_k|^2}{|\mathrm{e}^{\mathrm{i}\theta} - a_k|^2}. \tag{27}$$

Theorem 7.1 (Stop suppression forces dwell time). *Let $I \subset \mathbb{T}$ be a connected arc of angular length $L > 0$. If $\sup_I |p| \le \varepsilon < 1$, then*

$$\sup_{\theta \in [-\pi, \pi]} \tau_B(\theta) \ge \frac{LS}{16 \arcsin \varepsilon}. \tag{28}$$

Proof. The condition $|1 + B| \le 2\varepsilon$ confines the phase of B to an arc of half-width $2 \arcsin \varepsilon$ around π. Because I is connected and $\varepsilon < 1$, a continuous unwrapped phase cannot change branches inside I; its total variation is at most $4 \arcsin \varepsilon$. On the other hand, each Poisson kernel in (27) obeys the uniform lower bound

$$\frac{1 - |a_k|^2}{|\mathrm{e}^{\mathrm{i}\theta} - a_k|^2} \ge \frac{1 - |a_k|}{4}.$$

Integration over I therefore yields

$$\frac{L}{4} \sum_{k=1}^{S} (1 - |a_k|) \le 4 \arcsin \varepsilon.$$

For some k_*, $1 - |a_{k_*}| \le 16 \arcsin(\varepsilon)/(LS)$. At $\theta = \arg a_{k_*}$, its Poisson kernel is $(1 + |a_{k_*}|)/(1 - |a_{k_*}|)$, which alone is at least the right-hand side of (28). ☐

Corollary 7.2 (Delay-budget frontier). *Under the hypotheses of Theorem 7.1, let $D = \sup_\theta \tau_B(\theta)$. Then*

$$\varepsilon \ge \sin\left(\frac{LS}{16D}\right). \tag{29}$$

Consequently, if $\delta = \varepsilon^2$ is the corresponding rate envelope, the rate-improvement factor satisfies

$$\frac{1}{\delta} \le \csc^2\left(\frac{LS}{16D}\right) \sim \left(\frac{16D}{LS}\right)^2 \qquad (D/S \to \infty). \tag{30}$$

In particular, exponential rate improvement requires exponential peak delay when L is fixed and S grows subexponentially relative to the target.

Proof. Rearrange (28) to obtain $\arcsin \varepsilon \ge LS/(16D)$ and use monotonicity of sine on $[0, \pi/2]$. Squaring and inverting gives (30); the asymptotic form follows from $\sin x \sim x$. ☐

For the repeated real zero in (2), the exact peak is

$$\tau_{B_S}^{\max} = S \frac{1 + r_S}{1 - r_S} = S(2\mathrm{e}^{\alpha S} - 1). \tag{31}$$

Thus both the general lower bound and the construction are exponential when the leakage is exponential, though the simple repeated-zero design is not claimed constant-optimal. In physical language, high-Q passive memory replaces active pumping. That is a real architectural benefit in settings where storage time is cheap and drive power is expensive, but not a removal of resource cost. Theorem 7.2 makes the trade explicit: the squared spectral-rate advantage can grow at most quadratically with available peak delay, up to the order and arc-length factors shown.

8 Numerical and exact verification

The proofs above are analytic. Computation audits implementation identities, finite-precision behavior and scale. The public replay performs:

1. rational Gaussian elimination for Theorem 3.1;
2. closed-form pass-node generation and all $\binom{M}{3}$ Vandermonde minor checks;
3. dense global contractivity and stopband sweeps;
4. direct 6×6 unitarity tests on 1001 boundary frequencies;
5. the rational root-count inequality; and
6. multiprecision evaluation of γ_S, χ_S, ν_S and both sides of the robustness window; and
7. the squared-rate, vacuum-covariance and group-delay identities.

No optimizer is used to establish the theorem.

For the audit we fix $\alpha = 0.45$. Table 1 reports the observed stop norm ρ_S, its square, the reciprocal rate factor, the exact peak delay, and the maximum six-port unitarity residual. The smallest full-spark minor remains positive but becomes poorly conditioned by $S = 19$; this is why the numerical study stops there rather than presenting meaningless larger orders.

S	M	ρ_S	ρ_S^2	$1/\rho_S^2$	$\tau_{\max}$
5	15	$1.269\,10^{-1}$	$1.610\,10^{-2}$	$6.21\,10^{1}$	$8.99\,10^{1}$
7	21	$6.245\,10^{-2}$	$3.900\,10^{-3}$	$2.56\,10^{2}$	$3.20\,10^{2}$
9	27	$2.990\,10^{-2}$	$8.939\,10^{-4}$	$1.12\,10^{3}$	$1.02\,10^{3}$
11	33	$1.402\,10^{-2}$	$1.965\,10^{-4}$	$5.09\,10^{3}$	$3.09\,10^{3}$
15	45	$2.940\,10^{-3}$	$8.643\,10^{-6}$	$1.16\,10^{5}$	$2.56\,10^{4}$
19	57	$5.895\,10^{-4}$	$3.475\,10^{-7}$	$2.88\,10^{6}$	$1.96\,10^{5}$

Table 1: Dense-grid audit at $\alpha = 0.45$. All maximum calibration residuals are below 6.3×10^{-12} and all six-port unitarity errors below 6.0×10^{-14}. The theorem uses only $3(S-1)$ nodes; the replay audits all $3S$ interior nodes found for these orders.

The observed rate factor exceeds 2.8×10^{6} at $S = 19$, but that number should be read beside a peak delay near 2.0×10^{5} samples and a minimum Vandermonde minor near 1.3×10^{-14}. The result is a theorem about architecture and asymptotic possibility, not a claim that the largest audited device is robustly fabricable.

8.1 Passive-preserving tolerance stress test

To distinguish algebraic existence from component tolerance, we run a frozen Monte Carlo stress test at $S = 11$. Each of 2000 trials independently perturbs the logarithm of each pole gap $1 - r_g$ by a Gaussian standard deviation of 1%, each arm phase by 1% of the nominal pole gap, and each balanced interferometer angle by 10^{-3} radians. The seed is 20260810. All perturbations preserve $|B_g| = 1$ and use orthogonal beam splitters, so every trial remains a passive lossless six-port network; only its nominal calibration is lost.

The median stop norm is 0.01421, compared with the nominal 0.01402; its 95th percentile is 0.01474 and maximum is 0.01611. In contrast, the maximum residual over the *nominal* pass nodes has median 0.0700 and 95th percentile 0.1293. Thus stop suppression is locally stable in this component model, while exact high-order calibration is sensitive to pole placement. Retuning the pass frequencies restores exact roots for each lossless sample, but a fixed-frequency application must budget this error. The stress test supports the resource/conditioning caveat; it is not evidence for a robust comparator separation.

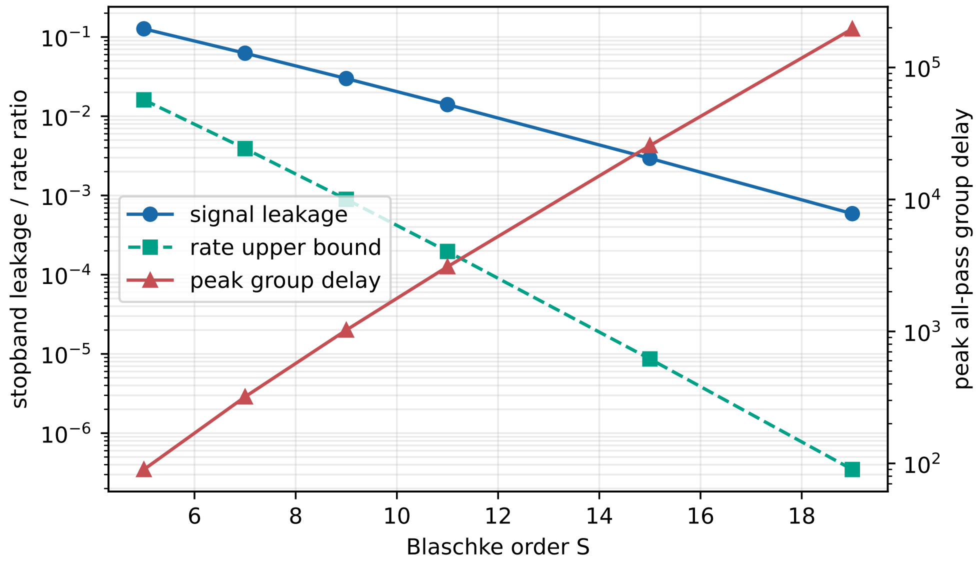

Figure 2: Signal leakage and its squared rate envelope fall exponentially, while peak all-pass group delay rises exponentially. Lines connect audited orders and are not statistical fits.

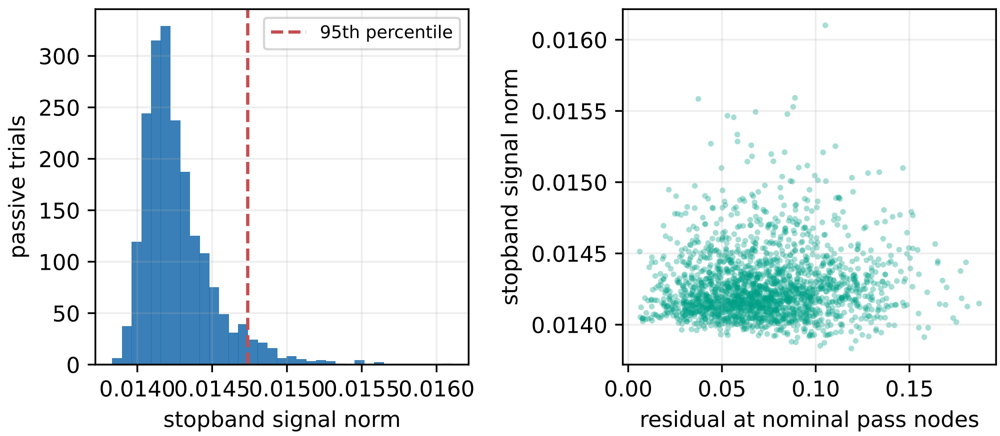

Figure 3: Passive-preserving tolerance audit at $S = 11$ (2000 frozen trials). Stop leakage changes mildly, whereas residuals at the nominal calibration nodes expose the high-Q phase sensitivity.

8.2 Multiprecision robustness certificate

Theorem 5.3 is analytic; its constants are finite because the three displayed matrices have full column rank. Nevertheless their scale is scientifically relevant. Table 2 evaluates them with 60 decimal digits and reports both the proved C_S and the necessary lower bound from Theorem 5.4. The enormous gap shows that our positive constant is not sharp. The rapid growth on both sides shows, without relying on Monte Carlo perturbations, that fixed absolute tolerance is the wrong asymptotic regime for these clustered nodes.

S	γ_S	χ_S	ν_S	$\log_{10} C_S$	necessary $\log_{10} C'_S$
5	$1.257\,10^{-3}$	$4.676\,10^{-16}$	$3.948\,10^{-5}$	23.873	1.105
7	$8.558\,10^{-5}$	$6.528\,10^{-21}$	$6.431\,10^{-7}$	31.825	1.584
9	$7.875\,10^{-6}$	$2.124\,10^{-28}$	$1.058\,10^{-8}$	42.235	2.019

Table 2: Multiprecision conditioning audit for the quantitative obstruction. The upper constant is a rigorous formula evaluated numerically, not a claim of optimality. The last column is a theorem-level lower bound on every possible constant in an inequality of the same form.

8.3 Observable rate gap above a vacuum floor

The six-port covariance identity in Theorem 6.2 is audited on a 257-point full-circle grid for every order in Table 1; the maximum operator-norm residual is 1.46×10^{-14}. Figure 4 plots (26) for $j_- = j_+ = 1$ using the observed leakage. With no baseline it reproduces the squared spectral gap; for every nonzero baseline it visibly approaches the predicted ceiling. For example, at $S = 9$ and $\kappa_0/j = 10^{-3}$ the certified total-rate factor is 528.5, while the zero-baseline excess-rate factor is 1118.7.

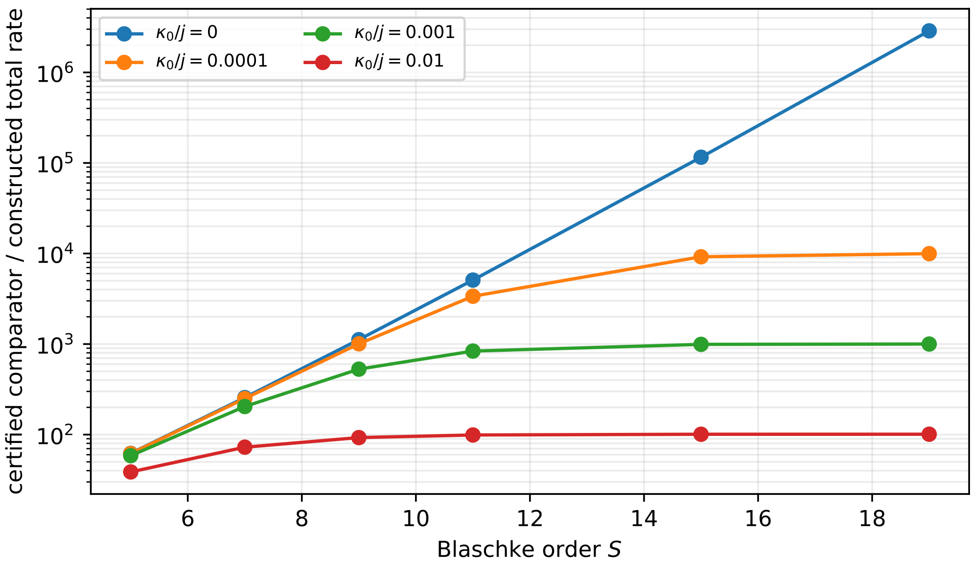

Figure 4: Certified total Ramsey-rate advantage after all vacuum inputs are included. The auxiliary ports produce a common, independently measurable floor rather than an uncontrolled architecture-dependent correction.

9 Relation to prior work and novelty boundary

Filter-function theory already supplies frequency-domain descriptions of controlled open-system error [1, 2], and experimental noise filtering is established [13]. Passive quantum feedback networks, lossless bounded-real transfers and cascade realizations are likewise established [3, 4, 5, 11]. Coherent passive reservoir engineering has, in particular, been formulated as a pole-placement route to decoherence-free subsystems [14]. Most directly, Techakesari and Nurdin preserve tangential responses and passivity in model reduction [6]; therefore neither "tangential quantum interpolation" nor "physical interpolation" is a novelty claim here. Tabak and Mabuchi connect Blaschke–Potapov factorization to trapped modes in delayed passive quantum networks [7]; therefore the use of Blaschke factors and the interpretation of near-boundary poles as trapped memory are also prior art.

The new mathematical claim is the conjunction proved in Theorems 5.1, 5.3, 5.4, 6.1, 6.2, 7.1 and 7.2: under one explicit collection of full-spark delayed calibrations and one fixed rational bidegree, a passive irreducible signal block has exponentially small stop leakage while every constant-channel reducing transfer retains unit leakage; finite defects obey an explicit conditioning bound that is provably nonuniform; a uniformly nondegenerate bath squares the excess-noise gap and vacuum completion yields a closed observable-rate formula; and strong scalar all-pass suppression forces a quantitative peak-delay cost. The earlier matrix-polynomial paper [10] supplied the classical root-count architecture mechanism. Here passivity, causal boundary phase, an explicit quantum-compatible loss completion, the Kossakowski transfer and the delay price change both the theorem and its interpretation. Delay–bandwidth and storage–bandwidth tradeoffs are classical themes; the paper claims neither theme in general. The candidate novelty is the explicit inequality used together with the fixed-bidegree architecture and bath-rate separations.

10 Limitations and falsification routes

Limitation 10.1 (Conditioned rather than uniform robustness). Theorem 5.3 propagates finite calibration and sampled reduction errors, but its certified constants are extremely conservative. Moreover, Theorem 5.4 proves that every theorem of this form deteriorates at least exponentially for the present clustered calibration family. A fabrication-oriented successor should redesign the nodes and signatures or optimize the finite interpolation constant; it cannot sharpen away the nonuniformity established here.

Limitation 10.2 (Linear stationary bath interface). Equations (20) and (24) assume a linear filter and a stationary Gaussian Markov bath. General non-Gaussian or nonlinear reservoirs need not transform by congruence, and strong-coupling or non-Markovian dynamics need not be summarized by a Davies or Ramsey rate. The vacuum ports are closed exactly inside the stated model, not outside it.

Limitation 10.3 (No many-body mass-gap theorem). The construction is a rational passive network, not an interacting quantum field theory. It neither proves a Yang–Mills mass gap nor establishes the interacting-buffer conjecture in [8]. Its relevance is methodological: it exhibits an exactly solvable architecture-dependent rate boundary and the memory resource behind it.

Limitation 10.4 (Restricted comparator). The no-go covers the static reciprocal split of Theorem 2.3 at the stated bidegree. A device with higher rational degree, a dynamical frequency-dependent change of channel coordinates, distinct input/output splits, feedback outside the modeled ports or active gain lies outside the theorem. Those alternatives spend or change resources and remain possible escape routes, not contradictions.

The most direct experimental falsification is a passive three-channel device that satisfies the same calibrated responses and bidegree but has a verified constant reducing line and stop norm below one. Such a device would contradict the algebraic root count and should be detectable from the public calibration ledger. A failure of the squared rate law under the stated linear-Gaussian assumptions would instead falsify the bath-transfer model.

11 Conclusion

Irreducible channel mixing is not merely a coordinate choice when calibration directions vary with frequency. An explicit passive six-port family exploits that fact to produce exponentially small signal leakage at a rational complexity where every constant-channel reducing comparator retains a unit branch. The obstruction survives finite calibration and approximate reduction errors with an explicit constant, while a scalar competitor proves that uniform tolerance is impossible for the chosen nodes. The bath spectral density squares the separation in excess dissipative rates; a closed Ramsey model turns every vacuum port into a measured common floor. The same construction exposes its own bill: exponentially long group delay, hence high-Q passive memory. Operationally, the Heisenberg cut can be displaced by architecture, but the resource boundary reappears as dwell time, vacuum noise and conditioning. That simultaneous positive result, robustness barrier and resource accounting is the central claim.

Code and data availability

Source, exact-arithmetic checks, dense and multiprecision audits, the JSON ledgers and the scripts that regenerate Figs. 1 to 4 are available in the public repository [9]:

`Public source and verification branch`

The principal commands are

```
python research/quantum_reservoir_filter.py
python research/quantum_reservoir_robustness.py
python research/quantum_reservoir_robust_bound.py
python research/quantum_reservoir_closed_dephasing.py
python verification/verify_passive_quantum_filter.py
```

The release record pins the PDF and result-ledger hashes. No private data, remote service or stochastic seed is required.

A A sharper view of the root-count obstruction

Suppose a reducing comparator has a subspace of dimension two rather than a declared reducing line. Its orthogonal complement is a reducing line because both the subspace and its complement are invariant by definition. Therefore the proof of Theorem 5.1 loses at most two calibrations, regardless of which nontrivial reducing dimension is initially supplied.

The degree $S+4$ is also conservative. If the scalar branch numerator has lower degree due to cancellation, the root-count contradiction becomes stronger. Boundary roots are legitimate because the denominator is assumed regular on $\mathbb{T}$. The proof does not require the comparator to be stable, causal, passive or quantum realizable; those restrictions only shrink its class.

B Bath congruence in a pure-dephasing model

For completeness, consider coupling operators $A_a = A_a^*$ and bath fields $B_a = B_a^*$ with interaction $H_I = \lambda \sum_{a=1}^{3} A_a \otimes B_a$. At a fixed secular frequency, collect the bath correlations in $J_{ab}(\omega)$. A linear prefilter with annihilation-amplitude transfer T maps the field vector to TB (with the loss ports initially in vacuum), so the signal contribution to its correlation matrix is TJT^*. For a transition

$m \leftrightarrow n$, let $c_a = \langle m|A_a|n\rangle$. The associated quadratic dissipative coefficient is, up to the common convention-dependent prefactor, c^*TJT^*c. This gives (20) for the excess-noise contribution. For the closed white-noise model, collect the vacuum covariance of all six inputs before adding J on the signal ports. The signal row identity $TT^* + LL^* = I$ then gives (24); no loss-port term remains unspecified. A colored or thermally populated auxiliary input would add its known term $LJ_{\mathrm{aux}}L^*$ and would no longer be architecture independent.

For commuting pointer observables, the same calculation applies at zero Bohr frequency to contrasts between pointer states. If the contrast vector is c, the off-diagonal density-matrix element decays as $\exp[-\kappa(c)t]$ with $\kappa(c)$ proportional to $c^*\Gamma_T(0)c$. This is the exact pure-dephasing setting in which the resource-boundary interpretation is cleanest.

C Reproducibility ledger

The JSON record stores, for each S:

- r_S, node count, numerator degree and root-count margin;
- minimum determinant magnitude over all signature triples;
- maximum calibration residual;
- global and stopband signal norms;
- the squared rate envelope and reciprocal factor;
- exact peak group delay; and
- maximum six-port unitarity residual.

Separate ledgers store the 60-digit singular margins and robustness constants, and the closed-vacuum covariance residual and total-rate frontiers. The analytic acceptance gates are deliberately looser than floating-point roundoff: calibration residual below 10^{-9}, unitarity residual below 10^{-10}, global norm below $1 + 2 \times 10^{-12}$, positive full-spark minors, and stop norm below (10). A failed gate exits nonzero.

References

[1] Gerardo A. Paz-Silva and Lorenza Viola. General transfer-function approach to noise filtering in open-loop quantum control. *Physical Review Letters*, 113:250501, 2014.

[2] Pascal Cerfontaine, Tobias Hangleiter, and Hendrik Bluhm. Filter functions for quantum processes under correlated noise. *Physical Review Letters*, 127:170403, 2021.

[3] John Gough and Matthew R. James. The series product and its application to quantum feedforward and feedback networks. *IEEE Transactions on Automatic Control*, 54(11):2530–2544, 2009.

[4] John E. Gough, Rolf Gohm, and Masahiro Yanagisawa. Linear quantum feedback networks. *Physical Review A*, 80:062104, 2009.

[5] John E. Gough and Guofeng Zhang. On realization theory of quantum linear systems. *Automatica*, 59:139–151, 2015.

[6] Onvaree Techakesari and Hendra I. Nurdin. Tangential interpolatory projection for model reduction of linear quantum stochastic systems. *IEEE Transactions on Automatic Control*, 62(1):5–17, 2017.

[7] Gil Tabak and Hideo Mabuchi. Trapped modes in linear quantum stochastic networks with delays. *EPJ Quantum Technology*, 3:3, 2016.

[8] Lluis Eriksson. The heisenberg cut as a resource boundary: An operational outlook from coherence maintenance costs. *ai.viXra.org*, 2512.0064v2, 2026.

[9] Lluis Eriksson. Passive quantum reservoir filter: Source, exact certificates, and numerical replays. GitHub repository, 2026. Public source and verification branch.

[10] Lluis Eriksson. Irreducible channel mixing and exponential calibration laws for matrix-polynomial and linear-phase MIMO FIR filters. GitHub release, 2026. Immutable source release at `https://github.com/lluiseriksson/finite-sample-spectral-certificates/releases/tag/v1.0-ai-vixra`.

[11] Hendra I. Nurdin, Stefanos Grivopoulos, and Ian R. Petersen. The transfer function of generic linear quantum stochastic systems has a pure cascade realization. *Automatica*, 69:324–333, 2016.

[12] E. B. Davies. Markovian master equations. *Communications in Mathematical Physics*, 39:91–110, 1974.

[13] A. Soare, H. Ball, D. Hayes, J. Sastrawan, M. C. Jarratt, J. J. McLoughlin, X. Zhen, T. J. Green, and M. J. Biercuk. Experimental noise filtering by quantum control. *Nature Physics*, 10:825–829, 2014.

[14] T. Nguyen, Z. Miao, Y. Pan, N. H. Amini, V. Ugrinovskii, and M. R. James. Pole placement approach to coherent passive reservoir engineering for storing quantum information. *IEEE Transactions on Automatic Control*, 63(9):2979–2993, 2018.