Irreducible Channel Mixing and Exponential Calibration Laws for Matrix-Polynomial and Linear-Phase MIMO FIR Filters

Lluis Eriksson

10 August 2026

Abstract

Scalar Chebyshev filtering treats every vector in a block Krylov iterate with the same polynomial, while simultaneously diagonalizable matrix coefficients amount to independent scalar filters after a fixed channel rotation. We study the larger class of Hermitian matrix-polynomial filters under tangential pass constraints. Already at fixed channel dimension $d = 3$, we construct rationally generated signatures with quantitative full-spark margin for which an irreducible symmetric filter has stopband leakage $O(e^{-cN})$. In contrast, every exactly calibrated symmetric coefficient family with any common nontrivial invariant channel subspace has leakage at least one; this comparator strictly contains the pairwise-commuting class and permits a fully noncommutative 2 by 2 block. A quantitative theorem covers approximately reducible filters by charging their sampled off-block coupling together with calibration error. The signature margin and the admissible combined error are both $e^{-O(N)}$, not e^{-N^3}. This exponential scale is unavoidable: for arbitrary pass nodes and signatures in the same separated bands, an explicit scalar binomial-tail polynomial has both pass error and stopband leakage at most $e^{-2N/81}$. Thus constant-error separation is impossible, while the irreducible construction achieves the correct exponential scale class. An affine cosine substitution gives the same robust law for fixed-latency reciprocal linear-phase MIMO FIR filters. We also give an exact rational five-tap certificate with stopband norm at most 25/32 and prove that its advantage survives reducible calibration errors $\delta < 7/1920$. A second exact-arithmetic certificate is calibrated from a public triaxial pump-vibration data set: it gives leakage below 0.96 versus one for every exactly calibrated reducible symmetric class, while its maximum directional error on six held-out records is 0.00385. In a frequency-domain-decomposition test on the frozen held-out cospectra, filtering rotates the leading modal direction by at most 0.222° and changes its leading spectral ordinate by at most 2.4×10^{-5} relatively. Numerical programs test the mechanism and expose an ordinary graph-denoising setting in which scalar Chebyshev filtering is instead preferable. Tangential matrix interpolation, MIMO filtering, scalar two-band approximation and convex FIR design are not claimed as new; the contribution is the fixed-order reducible/irreducible separation together with its two-sided calibration scale.

1 Problem and contribution

Let

$$P(x) = \sum_{k=0}^{N} C_k x^k, \qquad C_k \in \mathbb{R}^{d\times d},$$

and let $I_- = [-1, 0]$ be a stopband. Given distinct pass points $\lambda_j > 0$ and nonzero channel signatures $v_j \in \mathbb{R}^d$, impose

$$P(\lambda_j)v_j = v_j, \qquad 1 \le j \le M. \tag{1}$$

The leakage is

$$\rho(P) = \sup_{x\in I_-} \|P(x)\|_{\mathrm{op}}. \tag{2}$$

The coefficients will always be real symmetric. They need not commute.

The principal comparator is now the larger *reducible* subclass: the coefficient matrices have a common nonzero proper invariant channel subspace. In three dimensions, symmetry makes this equivalent to having a common invariant line, possibly beside a fully coupled, noncommuting 2 by 2 block. Pairwise-commuting coefficients form a strict subfamily: they admit a complete common orthogonal eigenbasis and reduce to scalar filters after a fixed channel rotation. The question is whether genuinely irreducible channel mixing can reduce (2) under the same pass constraints, and at what combined calibration/reducibility-error scale that comparison remains meaningful. For unit signatures define

$$\gamma(v_1, \ldots, v_M) = \min_{|J|=3} \sigma_{\min}\big([v_j]_{j \in J}\big).$$

For a unit direction u and a real-symmetric polynomial Q, also define its sampled off-block coupling

$$\beta_u(Q) = \max_j \|(I - uu^\top)Q(\lambda_j)u\|_2. \tag{3}$$

It vanishes when the coefficients share the invariant line $\mathbb{R}u$; conversely, because there are more pass nodes than the degree, vanishing at all pass nodes forces that common invariant line. Thus β_u is a directly evaluable relaxation of reducibility, rather than a commutator surrogate.

Our answer is an exponential calibration law. The input signatures are full spark and have an explicit $e^{-O(N)}$ lower margin. This makes the reducible obstruction quantitative under $e^{-O(N)}$ calibration and off-block errors, while a separate scalar construction proves that no asymptotic advantage can survive calibration errors bounded below independently of N.

$$\{\text{pairwise-commuting symmetric}\} \subsetneq \{\text{reducible symmetric}\} \subsetneq \{\text{all symmetric}\}.$$

The commuting class is a fixed orthogonal bank of scalar filters. The larger reducible class may contain a fully noncommuting 2 by 2 block but retains one common invariant channel line. Theorem 2.1 excludes this entire middle class; it does not claim optimality against other irreducible symmetric, nonsymmetric, rational, or state-space MIMO filters.

Figure 1: Comparator hierarchy and the exact scope of the separation.

2 Main separation theorem

Theorem 2.1 (Robust exponential tangential separation)**.** *There are absolute constants $C, c > 0$ and $N_0 \in \mathbb{N}$ such that for every $N \geq N_0$, with $d = 3$ and $M = N + 3$, one can choose distinct rational pass points $\lambda_1, \ldots, \lambda_M \in [1, 7/2]$ and full-spark unit vectors $v_1, \ldots, v_M \in \mathbb{R}^d$ satisfying*

$$\gamma(v_1, \ldots, v_M) \geq \underline{\gamma}_N := \frac{2^{-39N}}{48(N+4)^6}, \tag{4}$$

with the following properties.

1. *There is a degree-N polynomial P_N with real symmetric coefficient matrices satisfying* (1) *and*
$$\rho(P_N) \leq Ce^{-cN}.$$

2. *If a degree-N polynomial Q_N has real-symmetric coefficients with a common nonzero proper invariant subspace and satisfies the same tangential constraints, then $\rho(Q_N) \geq 1$. If its coefficients commute pairwise, then the stronger conclusion $Q_N(x) = I_d$ holds for every x.*

3. *More generally, let Q_N be any degree-N real-symmetric polynomial obeying*

$$\max_j \|(Q_N(\lambda_j) - I)v_j\|_2 \le \delta,$$

Then, for every unit $u \in \mathbb{R}^3$, with $L_N = (24e/5)^N$,

$$\rho(Q_N) \ge 1 - \frac{\sqrt{3}L_N}{\underline{\gamma}_N}\big(\delta + \beta_u(Q_N)\big). \tag{5}$$

Consequently $\rho(Q_N) \ge 3/4$ whenever $\delta + \beta_u(Q_N) \le \underline{\gamma}_N/(4\sqrt{3}L_N)$ for some u.

For every $N \ge N_0$, the coefficient matrices of P_N are irreducible over $\mathbb{R}$; in particular, they cannot commute pairwise.

The admissible combined error in the third item is $e^{-O(N)}$. We split the proof into an invariant-line obstruction, a small-leakage base construction, a quantitatively full-spark rational perturbation, and a matching scale obstruction.

2.1 The invariant-line obstruction

Lemma 2.2 (Exact invariant-line obstruction)**.** *Let $M \ge d + N$, let the λ_j be distinct, and let the $v_j \in \mathbb{R}^d$ be full spark. If $Q(x) = \sum_{k=0}^N D_k x^k$ has real-symmetric coefficients with a common unit eigenvector u and $Q(\lambda_j)v_j = v_j$ for every j, then $u^\top Q(x)u = 1$ for every x and $\rho(Q) \ge 1$ whenever 0 belongs to the stopband. If the coefficients commute pairwise, then $Q = I_d$.*

Proof. Write $q(x) = u^\top Q(x)u$. The common-eigenvector assumption and symmetry give $Q(x)u = q(x)u$. Taking the u component of the jth interpolation constraint gives

$$(q(\lambda_j) - 1)\langle u, v_j\rangle = 0.$$

A nonzero vector can be orthogonal to at most $d-1$ of the full-spark targets: otherwise d targets would lie in the $(d-1)$-dimensional hyperplane $u^\perp$ and could not span $\mathbb{R}^d$. Thus at least $M - d + 1 \ge N + 1$ distinct values of λ_j are roots of $q - 1$. Consequently $q = 1$, and $\|Q(0)\|_{\mathrm{op}} \ge |u^\top Q(0)u| = 1$.

If the coefficients commute pairwise, simultaneous orthogonal diagonalization supplies a common eigenbasis. Applying the preceding argument to every basis vector makes every diagonal scalar entry equal to one, hence $Q = I_d$. □

For completeness, $\beta_u(Q) = 0$ at the $M > N$ distinct pass nodes is not merely a sampled coincidence. For every $z \perp u$, the scalar polynomial $z^\top Q(x)u$ has degree at most N and vanishes at all those nodes, hence is identically zero. Therefore every coefficient preserves $\mathbb{R}u$. The root-count implication used here and in Lemma 2.2 is separately kernel-checked in Lean/mathlib, without user axioms or admitted goals, in the formal replay. The geometric full-spark count and the quantitative analytic estimates remain conventional proofs, not formalized claims.

The enlargement over the old comparator is strict. For example, coefficient matrices of the form

$$D_k = a_k \oplus B_k, \qquad B_0 = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, \quad B_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

share the first coordinate line but need not commute because $[B_0, B_1] \ne 0$. Lemma 2.2 still forces the scalar branch along that line to equal one, which alone gives the leakage lower bound.

2.2 A three-channel diagonal base filter

Put $M = N + 3$ and $S = \lceil M/3 \rceil$. Partition M nodes into three groups of sizes $s_g \in \{S-1, S\}$, $g = 0, 1, 2$, and define

$$\lambda_{g,r} = 1 + \frac{5(3r+g)}{6S}, \qquad 0 \le r < s_g. \tag{6}$$

The nodes are distinct rationals in $[1, 7/2]$, their global separation is at least $5/(6S)$, and consecutive nodes within one group are separated by $h = 5/(2S)$.

Write T_m for the Chebyshev polynomial of the first kind, set $\phi(x) = 2x + 1$, and let $L_{g,r}$ be the Lagrange cardinal polynomial for the nodes in group g. With $m_g = N - s_g + 1$, define

$$q_g(x) = \sum_{r=0}^{s_g-1} L_{g,r}(x) \frac{T_{m_g}(\phi(x))}{T_{m_g}(\phi(\lambda_{g,r}))}, \qquad P_0(x) = \operatorname{diag}(q_0(x), q_1(x), q_2(x)). \tag{7}$$

Each q_g has degree at most N and equals one at every node in its group.

Lemma 2.3. *Let*

$$\eta_0 = \operatorname{arcosh}(3), \quad \eta_1 = \operatorname{arcosh}(8), \quad \Delta\eta = \eta_1 - \eta_0, \quad K = \frac{36e}{5},$$

and set

$$c = \frac{2\eta_0 - \log K}{3} > 0, \qquad C_0 = 2\exp\left(\frac{2}{3}(\log K + \eta_0)\right).$$

For every $N \ge 7$,

$$\sup_{x \in I_-} \|P_0(x)\|_{\mathrm{op}} \le C_0 e^{-cN}.$$

Moreover, at every pass node,

$$\|P_0(\lambda_j)\|_{\mathrm{op}} \le B_N := 2(4e)^{S-1} e^{N\Delta\eta}.$$

Proof. For $x \in I_-$, $|x - \lambda_{g,r}| \le 9/2$. Equispacing within a group gives

$$\sum_{r=0}^{s_g-1} |L_{g,r}(x)| \le \frac{(18S/5)^{s_g-1}}{(s_g-1)!}.$$

Here we used $\sum_r [r!(s_g-1-r)!]^{-1} = 2^{s_g-1}/(s_g-1)!$. For $N \ge 7$, $S \ge 4$, $s_g \ge S-1$, and hence $S/(s_g-1) \le 2$. The bound $n! \ge (n/e)^n$ therefore gives

$$\sum_r |L_{g,r}(x)| \le K^{s_g-1}. \tag{8}$$

Also, for $\lambda \ge 1$,

$$T_m(\phi(\lambda)) \ge \tfrac{1}{2} e^{m\eta_0}.$$

Combining this with (8) yields

$$|q_g(x)| \le 2K^{s_g-1} e^{-(N-s_g+1)\eta_0}.$$

Since $s_g - 1 \le S - 1 \le (N+2)/3$, the claimed $C_0 e^{-cN}$ follows. The constant c is positive because $e < 3$ gives $K < 108/5 < e^{2\eta_0} = (3+\sqrt{8})^2$.

At another pass node the numerator distances in the Lagrange factors are at most $5/2$, so the same calculation gives $\sum_r |L_{g,r}(\lambda_j)| \le (4e)^{s_g-1}$. The Chebyshev ratio is at most $2e^{N\Delta\eta}$, proving the pass-node bound. The operator norm of a diagonal polynomial is its largest entry. $\square$

2.3 An explicit full-spark perturbation

Index all pass nodes by $j = 1, \dots, M$, let $g(j)$ be the coordinate group of node j, and set $t_j = j/(M+1)$. Define

$$w_j = (1, t_j, t_j^2)^\top, \quad \varepsilon_N = 2^{-13N}, \quad \widetilde{v}_j = e_{g(j)} + \varepsilon_N w_j, \quad v_j = \frac{\widetilde{v}_j}{\|\widetilde{v}_j\|_2}. \tag{9}$$

Lemma 2.4 (Quantitative rational full spark)**.** *The vectors in* (9) *are full spark and satisfy* (4) *for every $N \geq 7$.*

Proof. Fix any three indices and write

$$D(z) = \det(e_{g(j_1)} + zw_{j_1}, e_{g(j_2)} + zw_{j_2}, e_{g(j_3)} + zw_{j_3}) = \sum_{r=0}^{3} d_r z^r.$$

The cubic coefficient is the nonzero Vandermonde determinant of the three w_j. Every d_r is rational with denominator dividing $(M+1)^6$; therefore every nonzero coefficient has magnitude at least $(M+1)^{-6}$. Expanding the determinant into its six signed products also gives $|d_r| \leq 18$.

Let r_0 be the first nonzero coefficient. For $N \geq 7$, $\varepsilon_N \leq [108(M+1)^6]^{-1}$, and hence

$$\sum_{r>r_0} |d_r| \varepsilon_N^{r-r_0} \leq 18(\varepsilon_N + \varepsilon_N^2 + \varepsilon_N^3) \leq \frac{1}{2(M+1)^6}.$$

It follows that

$$|D(\varepsilon_N)| \geq \frac{\varepsilon_N^{r_0}}{2(M+1)^6} \geq \frac{\varepsilon_N^3}{2(M+1)^6}.$$

Each unnormalized column has norm below two, so normalization loses at most a factor eight in the determinant. For a matrix with three unit columns the product of its two largest singular values is at most three. Its smallest singular value is consequently at least

$$\frac{\varepsilon_N^3}{48(M+1)^6} = \frac{2^{-39N}}{48(N+4)^6}.$$

The elementary inequality on ε_N holds at $N = 7$ and its left-to- right ratio decreases thereafter. $\square$

Lemma 2.5 (Stable symmetric interpolation)**.** *For every $N \geq 7$, there is a real-symmetric coefficient polynomial P_N satisfying $P_N(\lambda_j)\widetilde{v}_j = \widetilde{v}_j$ and*

$$\sup_{x \in I_-} \|P_N(x) - P_0(x)\|_{\text{op}} \leq e^{-cN}.$$

Proof. Let $\mathcal{S}_N$ be the finite-dimensional space of degree-at-most-N polynomials with real symmetric coefficients, and define

$$L_\varepsilon : \mathcal{S}_N \longrightarrow (\mathbb{R}^d)^M, \qquad L_\varepsilon(P) = \big(P(\lambda_j)(e_{g(j)} + \varepsilon w_j)\big)_{j=1}^M.$$

At $\varepsilon = 0$, this map is onto. Indeed, for an off-diagonal entry $p_{rs} = p_{sr}$, arbitrary output data prescribe values at the union of nodes assigned to coordinates r and s. That union contains at most $2S \leq N+1$ nodes for $N \geq 7$. Ordinary degree-N Lagrange interpolation is therefore onto on those nodes. A diagonal entry sees at most S nodes.

Equip the output space with the maximum Euclidean block norm and $\mathcal{S}_N$ with

$$\|P\|_X = \max\left\{ \sup_{x \in I_-} \|P(x)\|_{\text{op}}, \max_j \|P(\lambda_j)\|_{\text{op}} \right\}.$$

This also gives an explicit linear right inverse. To record its norm, let $y = (y_j)_{j=1}^{M}$ satisfy $\max_j \|y_j\|_2 \le 1$. For $r \ne g$, define the entry $(R_N y)_{rg} = (R_N y)_{gr}$ by scalar Lagrange interpolation on the union of groups r and g: its value is $(y_j)_r$ at a node in group g and $(y_j)_g$ at a node in group r. A diagonal entry uses only its own group. These prescriptions are consistent because the groups are disjoint, and the resulting symmetric polynomial obeys $L_0 R_N y = y$.

For any such scalar problem, write its $t \le 2S$ nodes in increasing order. Their separation $\delta = 5/(6S)$ implies that the denominator of the ith cardinal polynomial is at least

$$\delta^{t-1} i!(t-1-i)!.$$

On I_- and at every pass node, all numerator distances are at most $9/2$. Put $n = t-1$. Since every scalar problem contains at least $S-1$ nodes, $S/n \le 2$ for $S \ge 4$. Consequently the sum of the absolute cardinal functions is bounded by

$$\left(\frac{27S}{5}\right)^{t-1} \sum_{i=0}^{t-1} \frac{1}{i!(t-1-i)!} = \left(\frac{27S}{5}\right)^{t-1} \frac{2^{t-1}}{(t-1)!} \le \left(\frac{108e}{5}\right)^{n}.$$

The operator norm of a 3 by 3 matrix is at most three times its largest entry. Hence, with the norm $\|\cdot\|_X$ above,

$$\|R_N\|_{Y\to X} \le 6\left(\frac{108e}{5}\right)^{2S-1} =: \mathcal{R}_N. \tag{10}$$

Since $S = \lceil (N+3)/3 \rceil$, this explicit envelope satisfies $\mathcal{R}_N = \exp(O(N))$. Retaining the factorial is the step that improves the old superexponential perturbation scale.

Write $L_\varepsilon = L_0 + \varepsilon K_N$. Since $\|w_j\|_2 \le \sqrt{3}$, $\|K_N\|_{X\to Y} \le \sqrt{3}$. Hence

$$L_\varepsilon R_N = I + E_N, \qquad \|E_N\| \le \sqrt{3}\,\varepsilon \mathcal{R}_N.$$

For $\varepsilon = \varepsilon_N$ and $N \ge 7$, this is below $1/2$. Indeed, using $e < 11/4$, $e^{\Delta\eta} < 4$, $S \le (N+5)/3$, and $108e/5 < 64$ gives

$$\mathcal{R}_N \le 2^{4N+17}, \qquad B_N \le 2^{(10N+11)/3}. \tag{11}$$

Thus $\|E_N\| \le 2^{18-9N} < 1/2$, so $(I+E_N)^{-1}$ exists and has norm at most two.

Since $L_0(P_0) = (e_{g(j)})_j$, the residual at the perturbed targets is

$$r_N = (\widetilde{v}_j)_j - L_{\varepsilon_N}(P_0) = \varepsilon_N \big((I - P_0(\lambda_j))w_j\big)_j.$$

By Lemma 2.3,

$$\|r_N\|_Y \le \varepsilon_N \sqrt{3}(1 + B_N).$$

Set

$$\Delta P_N = R_N (I+E_N)^{-1} r_N, \qquad P_N = P_0 + \Delta P_N.$$

Then $L_{\varepsilon_N}(P_N) = (\widetilde{v}_j)_j$ and

$$\|\Delta P_N\|_X \le 2\sqrt{3}\,\mathcal{R}_N \varepsilon_N (1+B_N) \le e^{-cN}.$$

Indeed, $1 + B_N \le 2B_N$, so (11) bounds the correction by $2^{(71-17N)/3}$. Since $c < \log 2$, the ratio of this envelope to e^{-cN} is at most $2^{(71-14N)/3} \le 2^{-9}$ for $N \ge 7$. Every coefficient remains real symmetric because R_N maps into $\mathcal{S}_N$. Appendix A restates this construction as a four-object audit chain and records every numerical envelope used above. □

Lemma 2.6 (Quantitative near-reducibility obstruction)**.** *For the nodes and signatures above, let Q be any degree-N polynomial with real-symmetric coefficients and suppose*

$$\max_j \|(Q(\lambda_j) - I)v_j\|_2 \le \delta.$$

Then (5) *holds for every unit $u \in \mathbb{R}^3$.*

Proof. Fix a unit vector u and put $q(x) = u^\top Q(x)u$. At each pass node set $r_j = (I - uu^\top)Q(\lambda_j)u$. Symmetry and the orthogonal decomposition of v_j give

$$u^\top(Q(\lambda_j) - I)v_j = (q(\lambda_j) - 1)u^\top v_j + r_j^\top(I - uu^\top)v_j.$$

The left-hand side has magnitude at most δ, while $\|r_j\|_2 \le \beta_u(Q)$. If three signatures all had $|u^\top v_j| < \underline{\gamma}_N/\sqrt{3}$, then the Euclidean norm of the three overlaps would be below $\underline{\gamma}_N$, contradicting (4). Thus at most two signatures have such a small overlap, and at $N+1$ distinct nodes

$$|q(\lambda_j) - 1| \le \frac{\sqrt{3}(\delta + \beta_u(Q))}{\underline{\gamma}_N}.$$

For any selected $N+1$ nodes, ordered increasingly, their separation is at least $5/(6S)$. The sum of the absolute degree-N Lagrange cardinal values at zero is at most

$$\left(\frac{42S}{5}\right)^N \frac{1}{N!} \le \left(\frac{24e}{5}\right)^N = L_N,$$

where $S/N \le 4/7$ for $N \ge 7$. Hence $|q(0) - 1| \le \sqrt{3}L_N(\delta + \beta_u(Q))/\underline{\gamma}_N$. Finally $\|Q(0)\|_{\mathrm{op}} \ge |u^\top Q(0)u| = |q(0)|$, and $0 \in I_-$. $\square$

Proof of Theorem 2.1. The normalized and unnormalized interpolation constraints are equivalent. Combine Lemmas 2.3, 2.4, and 2.5. The first conclusion holds with $C = C_0 + 1$ and the $c > 0$ of Lemma 2.3. For the second conclusion, a common nontrivial invariant subspace of symmetric 3 by 3 matrices yields a common invariant line: use the subspace itself if it is one-dimensional and its invariant orthogonal complement otherwise. Lemma 2.2 then gives leakage at least one, and gives the identity conclusion in the commuting subcase. The third is Lemma 2.6. Choose $N_0 \ge 7$ so that $Ce^{-cN_0} < 3/4$. Then the constructed coefficients cannot commute pairwise or have any common nontrivial invariant subspace for $N \ge N_0$, and the robust comparison is strict at the displayed combined-error radius. $\square$

2.4 Why the exponential calibration scale is unavoidable

The preceding radius is small, but its exponential order cannot be replaced by a constant. The obstruction is already scalar and is independent of the signature geometry.

Proposition 2.7 (Universal scalar noise barrier)**.** *For every $N \ge 1$, arbitrary pass nodes $\lambda_j \in [1, 7/2]$, and arbitrary unit signatures v_j, there is a scalar degree-N polynomial s_N such that*

$$\sup_{x\in[-1,0]} |s_N(x)| \le e^{-2N/81}, \qquad \max_j |s_N(\lambda_j) - 1| \le e^{-2N/81}.$$

Consequently the commuting symmetric polynomial $s_N I$ has both leakage and maximum tangential calibration residual at most $e^{-2N/81}$.

Proof. Put $r(x) = 2(x+1)/9$ and $k_N = \lceil N/3 \rceil$. Define the binomial-tail polynomial

$$s_N(x) = \sum_{k=k_N}^{N} \binom{N}{k} r(x)^k (1 - r(x))^{N-k}.$$

On the stopband $0 \le r \le 2/9$, while on the pass interval $4/9 \le r \le 1$. Hoeffding's inequality for a binomial random variable [11] gives

$$\Pr\{X/N \ge 1/3\} \le e^{-2N/81} \quad (r \le 2/9)$$

and

$$\Pr\{X/N < 1/3\} \le e^{-2N/81} \quad (r \ge 4/9).$$

The ceiling in k_N only strengthens both deviations. These are precisely the two claimed bounds. $\square$

Thus exact feasibility yields a class separation, exponentially accurate feasibility retains it for the constructed signatures, and exponentially accurate scalar filters eventually erase it for every signature family. The sufficient and impossible tolerances do not have matching constants, but both have $-\log\delta = \Theta(N)$; an N-independent calibration floor is ruled out.

3 Block-Krylov meaning

The theorem has a direct interpretation for block polynomial filtering. Let $H = \sum_j \lambda_j q_j q_j^\top$ be real symmetric and let $B \in \mathbb{R}^{n\times d}$ be a starting block. For $P(x) = \sum_{k=0}^N C_k x^k$, define the right-coefficient block filter

$$\mathfrak{F}_P(H,B) = \sum_{k=0}^{N} H^k B C_k. \tag{12}$$

Its jth spectral row is

$$q_j^\top \mathfrak{F}_P(H,B) = (q_j^\top B) P(\lambda_j). \tag{13}$$

Proposition 3.1 (Pass preservation and stop contraction)**.** *Let $\mathcal{T}$ and $\mathcal{S}$ index target and stop eigenpairs. Assume the rows $b_j^\top = q_j^\top B$ obey*

$$b_j^\top P(\lambda_j) = b_j^\top \quad (j \in \mathcal{T}), \qquad \lambda_j \in I_- \quad (j \in \mathcal{S}).$$

Then the target spectral component is preserved exactly, and

$$\|Q_{\mathcal{S}}^\top \mathfrak{F}_P(H,B)\|_F \le \rho(P) \|Q_{\mathcal{S}}^\top B\|_F,$$

where $Q_{\mathcal{S}}$ has columns q_j, $j \in \mathcal{S}$. If the C_k are pairwise-commuting symmetric matrices, one fixed orthogonal channel rotation reduces (12) *to d independent scalar polynomial filters.*

Proof. The pass statement follows from (13). For the stop part,

$$\|Q_{\mathcal{S}}^\top \mathfrak{F}_P(H,B)\|_F^2 = \sum_{j\in\mathcal{S}} \|b_j^\top P(\lambda_j)\|_2^2 \le \rho(P)^2 \sum_{j\in\mathcal{S}} \|b_j\|_2^2.$$

For the final claim, simultaneously diagonalize all $C_k = U\Lambda_k U^\top$. Then the columns of BU are acted on by the scalar polynomials whose coefficients are the diagonal entries of the Λ_k. □

Because every coefficient in Theorem 2.1 is symmetric, its column constraints transpose to the row constraints in Proposition 3.1. The theorem therefore gives block spectral instances with exponentially small stop contraction that no fixed-basis collection of scalar pass-preserving filters can reproduce. This is an approximation-theoretic separation; an end-to-end runtime advantage also depends on filter construction, recurrence stability, orthogonalization, and the cost model.

4 Reciprocal linear-phase MIMO FIR meaning

The polynomial separation is equivalently a fixed-order statement for a standard multichannel FIR class. Let

$$x(y) = \frac{9y+5}{4}, \qquad G_N(\omega) = P_N(x(\cos\omega)).$$

Because G_N is a real symmetric matrix cosine polynomial of degree N, it has a unique expansion

$$G_N(\omega) = A_N + 2\sum_{r=1}^{N} A_{N-r}\cos(r\omega), \qquad A_r = A_r^\top.$$

The causal response

$$H_N(e^{i\omega}) = \sum_{r=0}^{2N} A_r e^{-ir\omega} = e^{-iN\omega} G_N(\omega), \qquad A_r = A_{2N-r}, \tag{14}$$

therefore has $2N+1$ real symmetric palindromic taps and common group delay N.

Corollary 4.1 (Fixed-latency MIMO FIR separation)**.** *For every $N \ge N_0$ there are $M = N+3$ distinct frequencies $\omega_j \in [0, \arccos(-1/9)]$ and full-spark directions $v_j \in \mathbb{R}^3$ for which a causal response of the form* (14) *satisfies*

$$H_N(e^{i\omega_j})v_j = e^{-iN\omega_j}v_j, \qquad \sup_{\omega\in[\arccos(-5/9),\pi]} \|H_N(e^{i\omega})\|_{\mathrm{op}} \le Ce^{-cN}.$$

Every response with the same real-symmetric palindromic tap structure whose taps have a common nontrivial invariant channel subspace and which satisfies the same calibrations has high-frequency norm at least one. If the taps commute pairwise, it is the pure delay $e^{-iN\omega}I$. More generally, put $G(\omega) = e^{iN\omega}H(e^{i\omega})$ and

$$\beta_u^{\mathrm{FIR}}(H) = \max_j \|(I - uu^\top)G(\omega_j)u\|_2.$$

For any symmetric palindromic comparator with phase-corrected calibration residual at most δ, its high-frequency norm is at least the right-hand side of (5) *with $\beta_u(Q_N)$ replaced by $\beta_u^{\mathrm{FIR}}(H)$. Conversely, for arbitrary calibration directions and frequencies in the same pass range, there is a scalar palindromic response with both phase-corrected pass error and high-frequency norm at most $e^{-2N/81}$.*

Proof. Take the data of Theorem 2.1 and choose $\cos\omega_j = (4\lambda_j - 5)/9$. The pass interval $[1, 7/2]$ maps into $[-1/9, 1]$, while the stop interval $[-1, 0]$ maps onto $[-1, -5/9]$. Equation (14) gives the calibration and transfers the norm bound unchanged. The affine substitution is invertible on coefficient spaces, as is the change between powers of $\cos\omega$ and the cosine basis. Thus the transformed coefficient matrices commute pairwise exactly when the original polynomial coefficients do, and the invertible coefficient changes preserve common invariant subspaces. Palindromy makes the FIR taps an invertible relabelling and rescaling of the cosine coefficients. Applying Lemma 2.2 therefore forces a unit-gain scalar branch for every reducible response and the full identity in the commuting subcase. Lemma 2.6 transfers both residual and off-block bounds, and Proposition 2.7 transfers the universal scalar response. □

The reducible comparator in Corollary 4.1 is strictly broader than a bank of scalar linear-phase filters in one fixed basis: it also allows an arbitrary coupled, noncommuting 2 by 2 symmetric block beside the protected scalar branch. The result does not compare implementation cost, quantization robustness, rational/state-space filters, or nonsymmetric MIMO filters, and the asymptotic calibrated directions are constructed rather than taken from a deployed system.

5 Exact finite witnesses

5.1 A full-spark quadratic certificate

The asymptotic mechanism already occurs with rational data at degree two. Take

$$(\lambda_1, \lambda_2, \lambda_3, \lambda_4) = \left(\tfrac{1}{2}, 1, \tfrac{3}{2}, \tfrac{7}{2}\right), \qquad (v_1, v_2, v_3, v_4) = \left(\begin{pmatrix}1\\2\end{pmatrix}, \begin{pmatrix}0\\1\end{pmatrix}, \begin{pmatrix}-1\\2\end{pmatrix}, \begin{pmatrix}-1\\1\end{pmatrix}\right).$$

Every pair of targets is linearly independent, hence the tuple is full spark in $\mathbb{R}^2$. Set $P(x) = C_0 + xC_1 + x^2C_2$, with

$$C_0 = \begin{pmatrix} 61/96 & 7/48 \\ 7/48 & 61/96 \end{pmatrix}, \quad C_1 = \begin{pmatrix} 7/24 & 7/48 \\ 7/48 & 77/96 \end{pmatrix}, \quad C_2 = \begin{pmatrix} -7/24 & -7/24 \\ -7/24 & -7/16 \end{pmatrix}.$$

Direct substitution verifies $P(\lambda_j)v_j = v_j$ for all four targets, and

$$[C_0, C_1] = \begin{pmatrix} 0 & 343/4608 \\ -343/4608 & 0 \end{pmatrix} \neq 0.$$

Proposition 5.1 (Exact quadratic gap)**.** *The polynomial above obeys*

$$\sup_{x\in[-1,0]} \|P(x)\|_{\mathrm{op}} \leq \frac{25}{32} < 1.$$

Every pairwise-commuting real-symmetric quadratic polynomial satisfying the same four tangential constraints is identically I_2 and has leakage one.

Proof. For the upper bound, put $x = t - 1$ and express the leading principal minor and determinant of $(25/32)I_2 \pm P(t-1)$ in the Bernstein basis on $t \in [0,1]$. The leading-minor coefficients are

$$\left(\tfrac{35}{48}, \tfrac{7}{24}, \tfrac{7}{48}\right), \qquad \left(\tfrac{5}{6}, \tfrac{61}{48}, \tfrac{17}{12}\right),$$

and the determinant coefficients are, respectively,

$$\left(\tfrac{1421}{1536}, \tfrac{2597}{6144}, \tfrac{4655}{27648}, \tfrac{931}{18432}, 0\right),$$
$$\left(\tfrac{1}{16}, \tfrac{1711}{3072}, \tfrac{3835}{3456}, \tfrac{3707}{2304}, \tfrac{1525}{768}\right).$$

All are nonnegative and both leading minors are positive. Sylvester's criterion therefore gives $(25/32)I_2 \pm P(x) \succeq 0$ throughout the stopband, proving the norm bound. The zero final coefficient in the first determinant shows that the bound is attained at an endpoint rather than being a grid artifact.

For the commuting class, apply Lemma 2.2 with $M = 4 = d + N$. Every scalar common-eigenbasis entry has three distinct roots of $q - 1$ and is identically one. $\square$

The gap is stable under imperfect calibration. Write $\widehat{v}_j = v_j/\|v_j\|_2$.

Proposition 5.2 (Robust near-reducibility obstruction)**.** *Let Q be any real-symmetric quadratic polynomial satisfying*

$$\|Q(\lambda_j)\widehat{v}_j - \widehat{v}_j\|_2 \leq \delta, \qquad j = 1, \ldots, 4.$$

For every unit $u \in \mathbb{R}^2$, define $\beta_u(Q)$ by (3). Then

$$\sup_{x\in[-1,0]} \|Q(x)\|_{\mathrm{op}} \geq 1 - 60\big(\delta + \beta_u(Q)\big).$$

In particular, a reducible comparator has $\beta_u(Q) = 0$ for a common invariant direction and, if $\delta < 7/1920$, its leakage is strictly larger than the certified $25/32$ leakage of the irreducible witness.

Proof. For any two distinct normalized targets a, b in the tuple, $|a^\top b|^2 \leq 9/10$. Hence their two-column smallest singular value obeys

$$\sigma_{\min}^2 = 1 - |a^\top b| \geq \frac{1 - |a^\top b|^2}{2} \geq \frac{1}{20} > \frac{1}{25}.$$

For every unit vector u, at least one member of any pair therefore has $|u^\top \widehat{v}_j| \ge 1/(5\sqrt{2})$. Thus at most one target has smaller overlap. Fix u and put $q(x) = u^\top Q(x)u$. The same symmetric off-block decomposition used in Lemma 2.6 shows that at three distinct pass nodes

$$|q(\lambda_j) - 1| \le 5\sqrt{2}\big(\delta + \beta_u(Q)\big).$$

For each of the four possible three-node subsets of $\{1/2, 1, 3/2, 7/2\}$, the sum of the absolute Lagrange weights at $x = 0$ is at most eight. Hence $|q(0) - 1| < 60(\delta + \beta_u(Q))$, so $\|Q(0)\|_{\mathrm{op}} \ge 1 - 60(\delta + \beta_u(Q))$. Finally, $1 - 25/32 = 7/32$ gives the stated threshold. □

The exact-arithmetic replay `verification/verify_quadratic_filter_witness.py` checks all constraints, target minors, commutators, Bernstein coefficients, and the rational robust-calibration constants without a numerical solver.

5.2 An exact five-tap MIMO FIR certificate

The quadratic witness has an exact causal FIR realization. Substitute $x = (9\cos\omega + 5)/4$ and define the palindromic taps

$$B_0 = B_4 = \begin{pmatrix} -189/512 & -189/512 \\ -189/512 & -567/1024 \end{pmatrix},$$

$$B_1 = B_3 = \begin{pmatrix} -63/128 & -21/32 \\ -21/32 & -21/64 \end{pmatrix}, \quad B_2 = \begin{pmatrix} -149/768 & -665/768 \\ -665/768 & -235/1536 \end{pmatrix}.$$

Their response obeys the exact identity

$$\sum_{r=0}^{4} B_r e^{-ir\omega} = e^{-2i\omega}\big(B_2 + 2B_1\cos\omega + 2B_0\cos 2\omega\big) = e^{-2i\omega} P\left(\frac{9\cos\omega + 5}{4}\right).$$

Consequently it preserves the four directions, up to their common delay phase, at cosine values $-1/3, -1/9, 1/9, 1$ and has operator norm at most $25/32$ on $\omega \in [\arccos(-5/9), \pi]$. Every reducible symmetric palindromic five-tap response meeting the same four calibrations is the pure two-sample delay and has norm one. Moreover $[B_0, B_1] \ne 0$. All statements, including the tap conversion, are checked by the same exact-arithmetic replay. The same change of variables transfers Proposition 5.2: if a five-tap comparator obeys $\|e^{2i\omega_j} H(e^{i\omega_j})\widehat{v}_j - \widehat{v}_j\|_2 \le \delta$, its high-frequency norm is at least $1 - 60(\delta + \beta_u^{\mathrm{FIR}}(H))$ for every unit u.

5.3 An affine two-constraint precursor

For degree one and $d = 2$, take pass data

$$\lambda_1 = \tfrac{1}{2}, \quad v_1 = (1,0)^\top, \qquad \lambda_2 = 1, \quad v_2 = (1,1)^\top,$$

and $P(x) = A + xB$ with

$$A = \begin{pmatrix} 4/5 & 1/5 \\ 1/5 & 3/10 \end{pmatrix}, \qquad B = \begin{pmatrix} 2/5 & -2/5 \\ -2/5 & 9/10 \end{pmatrix}.$$

The constraints hold exactly and

$$[A, B] = \begin{pmatrix} 0 & -1/10 \\ 1/10 & 0 \end{pmatrix}.$$

Convexity of $x \mapsto \|A + xB\|_{\mathrm{op}}$ puts its maximum over I_- at an endpoint. Direct diagonalization gives

$$\rho(P) = \max\{\|A\|_{\mathrm{op}}, \|A - B\|_{\mathrm{op}}\} = \frac{1 + \sqrt{61}}{10} < 0.882.$$

If A and B commute, their common eigenbasis contains a direction with nonzero overlap with both nonorthogonal, nonparallel targets. Its scalar affine entry equals one at two points and hence everywhere. The commuting leakage is at least one.

6 Numerical stress test

The repository script `research/route_c_scaling_pilot.py` solves a grid SDP with symmetric coefficient variables. For degree 3, dimension 9, 12 perturbed targets and seed 7, the results are:

Perturbation	SDP objective	Validation maximum	Commutator norm
0.001	0.2983490	0.2983682	0.466925
0.010	0.3199793	0.3200281	0.391179
0.030	0.3977121	0.3977588	0.708604

The optimization uses 81 stopband points and validation uses 4001 independent points. Clarabel labels these solves `optimal_inaccurate`; they are diagnostics, not proof. The exact witnesses and Theorem 2.1 do not depend on solver output.

6.1 Controlled block-iteration benchmark

The script `research/route_c_block_krylov_benchmark.py` embeds the quadratic certificate in a 256-dimensional Hermitian eigenproblem. The four pass eigenvalues are $1/2, 1, 3/2, 2$, the other 252 eigenvalues fill $[-1, 0]$, and the initial stop/pass Frobenius ratio is five. We repeatedly apply the degree-two matrix filter and compare it with

$$q(x) = \frac{8x^2 + 8x + 1}{7},$$

the scalar degree-two Chebyshev stopband filter normalized only at $x = 1/2$. Let $\sin\theta$ denote the largest principal-angle sine from the filtered block to the clean two-dimensional block defined by the four prescribed pass rows. The scalar pass error below is minimized over one global rescaling, so it does not merely report scalar amplitude growth.

Iterations	Matrix stop/pass	Matrix $\sin\theta$	Scalar pass error	Scalar $\sin\theta$
0	5.0000	0.9918	0.0000	0.9918
1	2.2207	0.9638	0.7518	0.7847
5	0.3752	0.5622	0.9195	0.9802
10	0.0760	0.1385	0.9199	0.9819
20	0.00434	0.00810	0.9199	0.9820

The matrix pass error is below 1.2×10^{-10} through ten iterations and is 7.8×10^{-4} at iteration 20 because floating-point errors are amplified outside the preserved directions; exact arithmetic preserves them. The scalar filter has stopband maximum $1/7$ and initially suppresses contamination much faster, but its pass multipliers are $(1, 17/7, 31/7, 127/7)$; iteration consequently rotates the clean block away from the prescribed one. The fair commuting comparator under all four exact tangential constraints is the identity and remains at stop/pass ratio 5 and $\sin\theta = 0.9918$. This controlled construction demonstrates both the separation and its recurrence-conditioning limitation, not a runtime advantage on a natural problem.

6.2 A measured triaxial calibration and exact continuum certificate

We next use VBL-VA001, an open lab-scale pump-vibration data set with triaxial acceleration records sampled at 20 kHz [2, 12]. The task is frequency-domain-decomposition (FDD) preprocessing, not fault classification. FDD extracts mode-shape estimates as leading singular vectors of cross-spectral matrices at spectral peaks [6]; filter-induced perturbations are therefore a scientifically relevant failure mode, and filtering artifacts can corrupt modal estimates [21]. We selected twelve normal-condition records by the archive prefixes `normal_000` and `normal_001`, using all six records under the first prefix for fitting and retaining all six under the second prefix. The file names, CRC-32 and SHA-256 digests were frozen before optimization.

Each record was polyphase-decimated by eight, then divided into length-1024 Hann windows with 50% overlap. We averaged the real 3 by 3 cospectral matrix and, using the training split only, selected the five largest-trace peaks between 20 and 1000 Hz subject to an eight-bin separation. Their leading cospectral eigenvectors are the channel signatures. The frequency node is $x = \cos(2\pi f/2500)$. Table 1 reports the six-decimal rationalization used by the exact certificate; every displayed component and node differs from its measured value by less than 5×10^{-7}.

Lemma 6.1 (Modal-component invariance)**.** *Let the cospectral matrix at a modal frequency have the decomposition $G_j = \lambda_j v_j v_j^\top + R_j$, where $\lambda_j \geq 0$, $\|v_j\|_2 = 1$, and $R_j \succeq 0$. If a real-symmetric response P obeys $P(x_j)v_j = v_j$, then the reciprocal palindromic FIR response $H(e^{\mathrm{i}\omega}) = e^{-\mathrm{i}N\omega}P(\cos\omega)$ obeys*

$$H(e^{\mathrm{i}\omega_j})G_j H(e^{\mathrm{i}\omega_j})^* = \lambda_j v_j v_j^\top + P(x_j)R_j P(x_j).$$

Thus the rank-one modal component, its channel shape, and its amplitude are preserved exactly, up to one common N-sample delay, while the residual cospectrum is filtered.

Proof. The scalar phase cancels in HG_jH^*, and symmetry gives $P(x_j)v_j v_j^\top P(x_j) = v_j v_j^\top$. □

f (Hz)	x	v_x	v_y	v_z
100.0977	0.968522	−0.054531	0.996129	−0.068944
200.1953	0.876070	0.980644	0.167125	0.102012
744.6289	−0.296151	0.152303	0.986844	−0.054250
844.7266	−0.524590	0.869001	−0.350335	0.349432
944.8242	−0.720003	0.606651	0.716228	−0.344951

Table 1: Training frequencies and rationalized triaxial signatures from VBL-VA001.

Proposition 6.2 (Data-grounded rational certificate)**.** *For the five rational nodes and signatures in Table 1, there is a real-symmetric quadratic polynomial P satisfying $P(x_j)v_j = v_j$ exactly and*

$$[C_k, C_\ell] \neq 0 \quad \textit{for some } k,\ell, \qquad \sup_{-1 \leq x \leq -929/1000} \|P(x)\|_{\mathrm{op}} < \frac{24}{25}.$$

Every real-symmetric quadratic whose coefficients have a common nontrivial invariant subspace and which satisfies the same five calibrations has leakage at least one. In the pairwise-commuting subcase it is identically I_3.

Proof. All ten 3 by 3 target determinants are nonzero in exact rational arithmetic; the smallest normalized absolute determinant is $0.009489456048\ldots$. Lemma 2.2, with $d = 3$, $N = 2$ and $M = 5$, therefore proves the common-invariant-line conclusion. Any common nontrivial invariant subspace of real-symmetric 3 by 3 coefficients yields such a line by orthogonal complementation, proving the stated reducible-class bound and the stronger identity conclusion for commuting coefficients.

For the upper bound, fix the $(1,1)$, $(1,2)$ and $(2,2)$ entries of C_2 to -0.316171, -0.537470 and -0.909040, respectively, and solve the remaining 15 rational linear equations for the other symmetric entries. The exact replay verifies the calibrations and a nonzero commutator. On the 101-point rational grid of $[-1, -929/1000]$, Sylvester's criterion gives

$$(191/200)^2 I_3 - P(x_i)^2 \succ 0.$$

The same replay proves $\sup_{|x|\le 1}\|P'(x)\|_{\mathrm{op}} \le 5.740500963$. The grid spacing is $71/100000$, hence every point is within half a grid spacing of some x_i and

$$\|P(x)\|_{\mathrm{op}} < \frac{191}{200} + 5.740500963\frac{71}{200000} < 0.957038 < \frac{24}{25}.$$

Every comparison in this paragraph is performed on integers and rational fractions; the printed decimals only abbreviate upward-valid bounds. □

The held-out signatures lie between 0.129° and 0.460° from their training counterparts. Without refitting, the rational witness has maximum held-out residual 0.0038491 and dense-grid leakage 0.94870. The numerical comparators in Table 2 use the unrounded measured nodes and signatures. A general nonsymmetric MIMO polynomial attains much lower training leakage, but its held-out residual is 0.1045 and its largest coefficient Frobenius norm is 36.06, compared with 2.623 for the rational symmetric witness. This single split suggests a reciprocity–robustness tradeoff; it does not establish a statistical generalization theorem.

For a task-level replay, we froze the full 3 by 3 training and held-out cospectral matrices at the five peaks, applied $P(x_j)G_jP(x_j)$ on the held-out split, and recomputed the FDD leading eigenpair. Across the five frequencies, the angle between the unfiltered and filtered leading directions is at most 0.2217°; the leading eigenvalue ratio differs from one by at most 2.40×10^{-5}. These figures are outputs of the offline replay, not consequences of the exact certificate. Lemma 6.1 explains the specification: exact calibration preserves an ideal modal component, while the held-out calculation measures the effect of residual components and train–test drift.

Filter	Stop leakage	Fit tolerance	Held-out residual
Rational symmetric witness	< 0.96000	exact	0.003849
Symmetric SDP optimum	0.94867	exact	0.003849
General nonsymmetric SDP	0.14514	exact	0.10447
Scalar quadratic optimum	0.97520	0.01	0.01000
Best of 64 fixed bases	0.96966	0.01	0.01011
Exact reducible symmetric class	≥ 1	exact	0

Table 2: Measured-calibration controls. The 64-basis row is a seeded stress test, not a global optimization certificate over orthogonal bases.

Thus the measured instance closes both the purely constructed-data objection and the objection that tangential preservation has no task meaning: its calibration and validation use disjoint sensor records, FDD supplies the mode-shape objective, and the exact 0.96-versus-1 gap is not solver-dependent. It does not show improved fault diagnosis, damping estimation, hardware cost or deployment performance; those claims are outside the present experiment.

6.3 A failed graph-denoising application gate

We also tested the standard MIMO graph-filter form $Y = \sum_{k=0}^{2} S^k X C_k$ on two smooth fields over an 18 by 23 grid graph, using four low-frequency full-spark row signatures. The symmetric MIMO SDP collapsed to the identity (validated stop norm 0.99999992), while the one-point-normalized scalar Chebyshev filter had stop norm 1/17. After two iterations its best globally rescaled pass error was only 0.00524 and its desired-subspace sine was 0.0292, compared with 0.997 for the symmetric MIMO solution. A nonsymmetric MIMO fit eventually preserved the exact geometry, but was inferior at short iteration counts. Thus ordinary graph denoising does not justify the exact tangential specification, and no such application claim is made. We retain this negative result to prevent selecting an example merely because it matches the theorem's notation.

7 Relation to established work and scope of the claim

Tangential matrix-polynomial interpolation is classical [4, 7]. Structured matrix-polynomial interpolation with Hermitian, positive-semidefinite and related symmetries is also established [1]; that work treats full matrix values and lists the tangential structured variant as a future problem. Matrix-valued Nevanlinna–Pick theory handles supremum-norm analytic interpolation, including complexity constraints [5, 3]. Kuroiwa and Lindquist treat the bitangential problem with explicit McMillan-degree constraints [15]. Most directly, Stefanovski and Georgijević prove that real stable rational matrices can meet bitangential constraints while having arbitrarily small spectral norm on a proper frequency region [18]; their degree grows and they do not compare reducible and irreducible Hermitian polynomial coefficients. Thus small regional norm under matrix tangential constraints is not itself our contribution. Tangential interpolation is also a standard bridge between MIMO model reduction and Krylov projection [8]; modern block rational approximation includes matrix-valued weights, noisy-data comparisons, and tangential interpolation mechanisms [10]. Recent iterative MIMO schemes even select low-rank tangential data using maximum-error criteria [13]. Their maximum-error step selects the next interpolation frequency, whereas their proved monotonicity concerns a weighted H_2 error; monotone H_∞ bounds are stated as future work. Gallivan et al.'s main result instead gives a generically unique minimal-McMillan-degree rational interpolant. Neither result provides a fixed-degree Hermitian-polynomial minimax separation from the larger class with a common invariant channel subspace. Minimax polynomial filtering and SDP formulations have a long numerical-linear-algebra history [19], and matrix-polynomial spaces arise naturally in block Krylov methods [17]. General MIMO graph filters already use matrix taps in precisely the form $\sum_k S^k X C_k$ [9], while scalar FIR minimax design by convex optimization is classical [20]. Kootsookos already formulates fixed-length MIMO FIR approximation of matrix transfer functions in the uniform H_∞ norm, proves global error lower bounds, and gives MIMO and linear-phase design algorithms [14]. More closely, Ljungars and Fu optimize the operator-norm error of matrix-valued multi-channel linear-phase FIR filters on a frequency grid by SDP [16]. Their formulation has neither directional pass interpolation nor an invariant-subspace/reducibility comparator. Accordingly, fixed-length MIMO H_∞ approximation, MIMO linear phase, operator-norm minimax design, and SDP are not claimed as new.

To our knowledge, the candidate contribution is therefore narrower than norm-constrained or degree-constrained tangential interpolation or FIR design alone. Under fixed ordinary degree—equivalently fixed reciprocal FIR latency—a three-channel Hermitian-coefficient tangential minimax problem exhibits an exponential separation between irreducible coefficients and every reducible symmetric family. The excluded family is strictly larger than fixed-basis scalar banks: it permits a frequency-dependent, noncommuting 2 by 2 block, provided one channel line remains invariant. The quantitative result also identifies a joint calibration/reducibility scale: an exponentially small sum of tangential residual and sampled off-block coupling is sufficient for the constructed separation, whereas a universal scalar construction proves that an exponentially small calibration residual can erase every such separation. We have not found this two-sided law or the reducible/irreducible separation in the sources above.

The complete texts of two close restricted papers—Fuhrmann (2010) and Stefanovski–Georgijević (2016)—could not be obtained. Publisher portals, repository copies, metadata APIs, and independent network routes were checked. Their accessible statements already rule out broader novelty claims, but absence of a theorem from inaccessible text is not evidence of priority. The claim is therefore limited to the displayed theorem and to the literature actually inspected. The FDD lemma and held-out replay support a preprocessing interpretation; they do not constitute a deployment or diagnostic study.

Reproducibility. The source and research ledger are in the project repository. The repository links the three-channel theorem arithmetic replay, the exact rational replay, the block-iteration benchmark, its machine-readable output, the deliberately negative graph-filter pilot, the VBL-VA001 provenance manifest, its exact continuum replay, the offline baseline replay, and its machine-readable output, the clause-by-clause priority matrix, the Lean root-count proof, the single-command release replay, and the artifact manifest with the PDF hash.

AI-assistance disclosure. OpenAI Codex assisted with literature exploration, numerical and exact-arithmetic code, Lean formalization scaffolding, and drafting and editing the manuscript. The named author directed the research, selected the claims and scope, and accepts responsibility for the submitted content. No AI system is listed as an author.

A Expanded audit trail for stable symmetric interpolation

This appendix expands the quantitative chain in Lemma 2.5 so that each norm conversion and exponential envelope can be checked separately. It introduces no additional claim. Recall that $M = N + 3$, $S = \lceil (N+3)/3 \rceil$, $N \geq 7$, and $\varepsilon_N = 2^{-13N}$. The domain norm and data norm are

$$\|P\|_X = \max \left\{ \sup_{x \in I_-} \|P(x)\|_{\mathrm{op}}, \max_j \|P(\lambda_j)\|_{\mathrm{op}} \right\}, \qquad \|y\|_Y = \max_j \|y_j\|_2.$$

The four quantities that enter the correction are collected in Table 3.

Object	Definition	Explicit bound for $N \geq 7$
Right inverse	$L_0 R_N = I$ on $(\mathbb{R}^3)^M$	$\|R_N\|_{Y \to X} \leq \mathcal{R}_N \leq 2^{4N+17}$
Perturbation	$E_N = \varepsilon_N K_N R_N$	$\|E_N\| \leq 2^{18-9N} < 1/2$
Base residual	$r_N = \varepsilon_N((I - P_0(\lambda_j))w_j)_j$	$\|r_N\|_Y \leq \varepsilon_N \sqrt{3}(1 + B_N)$, $B_N \leq 2^{(10N+11)/3}$
Correction	$\Delta P_N = R_N(I + E_N)^{-1} r_N$	$\|\Delta P_N\|_X \leq 2^{(71-17N)/3} \leq e^{-cN}$

Table 3: The complete norm chain used to impose the perturbed calibrations.

1. Right-inverse geometry. Every diagonal scalar interpolation problem uses one coordinate group and every off-diagonal problem uses the union of two groups. Thus its node count t satisfies $S - 1 \leq t \leq 2S$ and its degree is at most $t - 1 \leq N$. Writing $n = t - 1$, the global node separation $5/(6S)$ and the numerator-distance bound $9/2$ give the exact cardinal sum

$$\sum_{i=0}^{t-1} |\ell_i(x)| \leq \left(\frac{27S}{5} \right)^n \sum_{i=0}^{n} \frac{1}{i!(n-i)!} = \left(\frac{27S}{5} \right)^n \frac{2^n}{n!}.$$

Because $S/n \leq 2$ and $n! \geq (n/e)^n$, this is at most $(108e/5)^n$. There are at most six independent symmetric entries, and $\|A\|_{\mathrm{op}} \leq 3 \max_{r,s} |A_{rs}|$ for a 3 by 3 matrix. The explicit entrywise construction in Lemma 2.5 therefore yields

$$\|R_N\|_{Y \to X} \leq 6 \left(\frac{108e}{5} \right)^{2S-1} = \mathcal{R}_N.$$

Using $108e/5 < 64$ and $S \leq (N+5)/3$ gives the deliberately loose but convenient power-of-two envelope $\mathcal{R}_N \leq 2^{4N+17}$. Keeping the factorial in the preceding display is what prevents an $e^{O(N \log N)}$ loss.

2. Invertibility after the full-spark perturbation. The map difference is $L_\varepsilon - L_0 = \varepsilon K_N$. Since every $w_j = (1, t_j, t_j^2)^\top$ has norm at most $\sqrt{3}$, $\|K_N\|_{X\to Y} \le \sqrt{3}$. Consequently

$$\|E_N\| = \|\varepsilon_N K_N R_N\| \le \sqrt{3}\, 2^{-13N} 2^{4N+17} \le 2^{18-9N} < \frac{1}{2}.$$

The Neumann series is therefore legitimate and $\|(I+E_N)^{-1}\| \le 2$. Moreover, $L_{\varepsilon_N} R_N = I + E_N$, so the displayed correction satisfies the calibrations algebraically rather than approximately:

$$L_{\varepsilon_N}(P_0 + \Delta P_N) = L_{\varepsilon_N} P_0 + (I+E_N)(I+E_N)^{-1} r_N = (\widetilde{v}_j)_j.$$

3. Residual and correction size. Lemma 2.3 gives $\max_j \|P_0(\lambda_j)\|_{\mathrm{op}} \le B_N$ and the elementary envelopes give $B_N \le 2^{(10N+11)/3}$. Hence

$$\|r_N\|_Y \le 2^{-13N}\sqrt{3}(1+B_N).$$

Combining this with the preceding two paragraphs yields

$$\begin{aligned}\|\Delta P_N\|_X &\le 2\sqrt{3}\,\mathcal{R}_N\, 2^{-13N}(1+B_N)\\ &\le 2^{(71-17N)/3}.\end{aligned}$$

Here the last line uses $2\sqrt{3} < 4$ and $1 + B_N \le 2B_N$. The base-filter exponent

$$c = \frac{2\,\mathrm{arcosh}(3) - \log(36e/5)}{3}$$

is positive and smaller than $\log 2$. Therefore

$$\frac{2^{(71-17N)/3}}{e^{-cN}} \le 2^{(71-14N)/3} \le 2^{-9} \qquad (N \ge 7),$$

which proves $\|\Delta P_N\|_X \le e^{-cN}$ with room to spare.

4. Reproduction boundary. The repository's constant-channel verification script replays the node arithmetic, determinant envelopes, Neumann bound, correction bound and scalar barrier without solving an optimization problem. The Lean file checks only the scalar root-count implication used by the invariant-line obstruction. Neither artifact replaces the conventional proof above; together they isolate the arithmetic and algebraic failure surfaces that are easiest to audit.

References

[1] Daniel Alpay and Izchak Lewkowicz. Interpolation by polynomials with symmetries. *Linear Algebra and its Applications*, 456:64–81, 2014.

[2] Bagus Tris Atmaja, Haris Ihsannur, Suyanto, and Dhany Arifianto. Lab-scale vibration analysis dataset and baseline methods for machinery fault diagnosis with machine learning. *Journal of Vibration Engineering & Technologies*, 12(2):1991–2001, 2024.

[3] Joseph A. Ball and Vladimir Bolotnikov. The bitangential matrix Nevanlinna–Pick interpolation problem revisited. *Linear Algebra and its Applications*, 516:122–188, 2017.

[4] Joseph A. Ball and Jeongook Kang. Matrix polynomial solutions of tangential Lagrange–Sylvester interpolation conditions of low McMillan degree. *Linear Algebra and its Applications*, 137–138:699–746, 1990.

[5] Anders Blomqvist, Anders Lindquist, and Ryozo Nagamune. Matrix-valued Nevanlinna–Pick interpolation with complexity constraint: an optimization approach. *IEEE Transactions on Automatic Control*, 48(12):2172–2190, 2003.

[6] Rune Brincker, Lingmi Zhang, and Palle Andersen. Modal identification of output-only systems using frequency domain decomposition. *Smart Materials and Structures*, 10(3):441–445, 2001.

[7] Paul A. Fuhrmann. On tangential matrix interpolation. *Linear Algebra and its Applications*, 433(11–12):2018–2059, 2010.

[8] Kyle A. Gallivan, An Vandendorpe, and Paul Van Dooren. Model reduction of MIMO systems via tangential interpolation. *SIAM Journal on Matrix Analysis and Applications*, 26(2):328–349, 2004.

[9] Fernando Gama, Antonio G. Marques, Alejandro Ribeiro, and Geert Leus. MIMO graph filters for convolutional neural networks, 2018.

[10] Ion Victor Gosea and Stefan Güttel. Algorithms for the rational approximation of matrix-valued functions. *SIAM Journal on Scientific Computing*, 43(5):A3033–A3054, 2021.

[11] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.

[12] Haris Ihsannur, Bagus Tris Atmaja, Suyanto, and Dhany Arifianto. VBL-VA001: Lab-scale vibration analysis dataset, 2022. Version 1.0, CC BY 4.0.

[13] Jared Jonas and Bassam Bamieh. An iterative tangential interpolation algorithm for model reduction of MIMO systems, 2026.

[14] Peter J. Kootsookos. *FIR(q) Filter Designs Using H_∞ Techniques*. PhD thesis, Australian National University, 1991.

[15] Yohei Kuroiwa and Anders Lindquist. Bi-tangential Nevanlinna–Pick interpolation with a complexity constraint. In *Proceedings of the 17th International Symposium on Mathematical Theory of Networks and Systems*, pages 2541–2560, Kyoto, Japan, 2006.

[16] Gustav Ljungars and Minyue Fu. Design of multi-channel linear phase FIR filters. In *Mini-Symposium on Mathematical Analysis and Its Applications*, pages 67–86, Newcastle, Australia, 1998.

[17] Michele Rinelli and Raf Vandebril. Block Krylov subspaces and orthogonal matrix polynomials: a structural correspondence with applications to unitary matrices, 2026.

[18] Jovan Stefanovski and Dušan Georgijević. Interpolation with constraint on frequency region and systems & control application. *Systems & Control Letters*, 97:70–82, 2016.

[19] Kim-Chuan Toh and Lloyd N. Trefethen. The Chebyshev polynomials of a matrix. *SIAM Journal on Matrix Analysis and Applications*, 20(2):400–419, 1998.

[20] Shih-Ping Wu, Stephen Boyd, and Lieven Vandenberghe. FIR filter design via spectral factorization and convex optimization. In Biswa N. Datta, editor, *Applied and Computational Control, Signals, and Circuits*, volume 1, pages 215–245. Birkhäuser, 1998.

[21] Zachariah Wynne, James R. Hopgood, Tim Stratford, and Thomas P. S. Reynolds. Reducing coherent filtering artefacts in time-domain operational modal analysis. *Structural Control and Health Monitoring*, 29(8):e2970, 2022.