Exact Memory of Direct-Sum Spectral Fan-Out: Generic Maximality, Colliding Targets, and the Three-Line Phase Diagram

Lluis Eriksson

11 August 2026

Abstract

At L boundary frequencies, a passive lossless network is required to send one fixed k-dimensional input channel to prescribed output k-planes. Mutually orthogonal targets make one exact fan-out law transparent. We show that orthogonality is inessential and that maximal memory is generic. If the L target planes are in direct-sum position, then every square finite rational-inner interpolant regular at the nodes has minimum McMillan degree

$$d_{\min} = k(L-1).$$

When $N \geq Lk$, this direct-sum locus is open and dense in $\mathrm{Gr}(k,N)^L$, which is the precise sense in which the maximum is generic. More generally, with r the dimension of their linear span,

$$r - k \leq d_{\min} \leq k(L-1).$$

The upper bound is constructive: all free diagonal blocks of the boundary Pick matrix are completed by one positive leading block and a zero Schur complement, producing a conservative realization with at most $k(L-1)$ states. The lower bound follows from a Pick–Stein displacement and is independent of frame gauges and node spacing.

The theorem exposes an exact-data discontinuity. We construct L direct-sum k-planes whose pairwise angles tend to zero. Their minimum degree remains $k(L-1)$ for every nonzero opening but falls to zero at collision; the certifying Gram margin is exactly $\epsilon^2/(1+\epsilon^2)$. A fail-closed noisy eigenvalue rule quantifies the precision needed to certify this cost. For three scalar target lines we also determine the complete phase diagram: degree zero for one line, degree one precisely for three distinct lines in a common two-plane, and degree two otherwise. Thus a repeated target can cost more than three distinct targets, a rigidity invisible to span alone. Deterministic code replays 512 arbitrary-node completions, 1,024 noisy certificates, 20 target-collision fixtures, four clustered-node fixtures, and five minimal unitary colligations, with an independent verifier and frozen machine-readable certificate.

Keywords: rational inner matrix; McMillan degree; boundary Nevanlinna–Pick interpolation; positive semidefinite completion; Stein displacement; spectral routing; passive network; Grassmannian.

1 The question and the answer

A finite passive lossless device is represented here by a square rational matrix S that is analytic on the unit disk and unitary almost everywhere on its boundary. Its McMillan degree counts the states of a minimal conservative realization. Fix an isometric input frame $X \in \mathbb{C}^{N\times k}$ and distinct nodes $\zeta_1, \dots, \zeta_L \in \mathbb{T}$. At node i only an output plane $\mathcal{Y}_i \in \mathrm{Gr}(k,N)$ is prescribed:

$$S(\zeta_i)\,\mathrm{ran}\,X = \mathcal{Y}_i, \qquad i = 1, \dots, L. \tag{1}$$

The interpolant is required to be regular at the nodes. Define

$$d_{\min}(\zeta, \mathcal{Y}; X) = \min\{\deg_{\mathrm{McM}} S : S \text{ is square rational inner and satisfies (1)}\}. \tag{2}$$

The value is independent of the particular input frame: constant unitaries move one k-plane to another without changing degree.

For $L = 1$, a constant unitary gives $d_{\min} = 0$. Henceforth $L \geq 2$. The mutually orthogonal example gives $d_{\min} = k(L-1)$ transparently, but does not decide whether orthogonality creates the cost or merely makes it visible. The answer is the latter.

Theorem 1.1 (Universal sandwich)**.** *Let $\zeta_1, \dots, \zeta_L$ be distinct points of $\mathbb{T}$ and let $\mathcal{Y}_1, \dots, \mathcal{Y}_L$ be k-planes in $\mathbb{C}^N$. Put*

$$r = \dim \operatorname{span}(\mathcal{Y}_1, \dots, \mathcal{Y}_L).$$

Then

$$r - k \leq d_{\min} \leq k(L-1). \tag{3}$$

Both inequalities are independent of node separation. The upper bound is certified by an explicit positive Pick completion.

Corollary 1.2 (Generic maximality)**.** *If the targets are in direct-sum position,*

$$\dim \operatorname{span}(\mathcal{Y}_1, \dots, \mathcal{Y}_L) = Lk,$$

then

$$d_{\min} = k(L-1). \tag{4}$$

When $N \geq Lk$, direct-sum position is open and dense in $\mathrm{Gr}(k, N)^L$. Hence the largest possible minimum degree is the generic value, not an exceptional orthogonal value.

Every pair drawn from a direct-sum table has minimum degree k, by applying Theorem 1.1 with $L = 2$. The ratio between the global minimum and the largest pairwise minimum is therefore $L - 1$. Pairwise feasibility does not aggregate into global economy.

Three features distinguish the present result from a restatement of the orthogonal construction. First, the target Gram matrix may be fully dense. Second, the upper bound now requires a block completion rather than a scalar Cauchy shift. Third, direct-sum targets may be arbitrarily close in the Grassmannian, producing the discontinuity proved in Section 5.

2 The Pick–Stein interface

Choose arbitrary isometric frames $W_i \in \mathbb{C}^{N \times k}$ with $\operatorname{ran} W_i = \mathcal{Y}_i$, and form the calibrated block Gram matrix

$$G = [W_i^* W_j]_{i,j=1}^L = W^* W, \qquad W = [W_1 \ \cdots \ W_L]. \tag{5}$$

Changing frames inside the planes conjugates G by a block-diagonal unitary. In particular,

$$\operatorname{rank} G = \operatorname{rank} W = r \tag{6}$$

is intrinsic.

Let

$$Z = \operatorname{diag}(\zeta_1, \dots, \zeta_L) \otimes I_k, \qquad J_L = \mathbf{1}\mathbf{1}^*.$$

The boundary Pick variables are the blocks of a Hermitian matrix $P = [P_{ij}] \in \mathbb{C}^{Lk \times Lk}$. Its off-diagonal blocks are fixed:

$$P_{ij} = \frac{I_k - W_i^* W_j}{1 - \overline{\zeta}_i \zeta_j}, \qquad i \neq j. \tag{7}$$

The diagonal blocks are free because both numerator and denominator vanish on the boundary. Equivalently, every completion satisfies the displacement identity

$$P - Z^* P Z = J_L \otimes I_k - G. \tag{8}$$

We isolate the standard interpolation bridge in the precise form used below.

Proposition 2.1 (Rank–degree bridge). *If a matrix $P \succeq 0$ satisfies* (8)*, then there is a square rational-inner S, regular at the prescribed nodes after removal of decoupled unit-circle modes, such that $S(\zeta_i)X = W_i$ and*

$$\deg_{\mathrm{McM}} S \leq \operatorname{rank} P.$$

Conversely, an interpolant of degree d produces such a positive P with $\operatorname{rank} P \leq d$ *after choosing the induced frames $W_i = S(\zeta_i)X$ in the target planes.*

Proof. Factor $P = F^*F$ with $F = [F_1 \ \cdots \ F_L]$ and $F_i \in \mathbb{C}^{q\times k}$, where $q = \operatorname{rank} P$. Identity (8) says that the two block-column families

$$\begin{bmatrix} \zeta_i F_i \\ X \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} F_i \\ W_i \end{bmatrix}$$

have the same Gram matrix. The induced isometry extends to a unitary colligation

$$U = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \quad \text{on } \mathbb{C}^q \oplus \mathbb{C}^N.$$

Its transfer matrix

$$S(z) = D + zC(I - zA)^{-1}B$$

is rational inner and has degree at most q. The interpolation equations give $(I - \zeta_i A)F_i = BX$ and then $S(\zeta_i)X = W_i$. If the chosen unitary extension contains a unit-circle state mode, contractivity forces that mode to be reducing and decoupled from B and C; deleting it leaves the transfer matrix and the interpolation data unchanged and makes all nodes regular.

Conversely, let S interpolate the prescribed planes and choose the induced frames $W_i = S(\zeta_i)X$. The de Branges–Rovnyak kernel of a finite rational-inner matrix,

$$K_S(z, w) = \frac{I - S(z)^*S(w)}{1 - \bar{z}w},$$

has a finite Kolmogorov factor of dimension $d = \deg_{\mathrm{McM}} S$. Taking regular boundary limits and compressing by X gives $P \succeq 0$ with $\operatorname{rank} P \leq d$, off-diagonal blocks (7), and hence (8). This is the finite-dimensional lurking-isometry form of matrix boundary Nevanlinna–Pick interpolation [5, 3, 7]. □

The proposition separates the new geometry from established realization machinery. Our contribution is not the Schur complement or the rank–degree bridge individually, but the sharp completion and lower bound that make them meet for generic subspace data.

3 A universal rank-$k(L-1)$ completion

Partition the partially specified matrix in (7) after its first $L-1$ block rows and columns. Let A_0 be the Hermitian $k(L-1) \times k(L-1)$ matrix whose off-diagonal blocks are the corresponding blocks in (7) and whose diagonal blocks are zero. Let B_0 stack the prescribed blocks $P_{1L}, \ldots, P_{L-1,L}$.

Lemma 3.1 (Zero-Schur completion). *For every $t > \|A_0\|$, set*

$$Q = tI + A_0, \qquad P = \begin{bmatrix} Q & B_0 \\ B_0^* & B_0^* Q^{-1} B_0 \end{bmatrix}. \tag{9}$$

Then $P \succeq 0$, it has all the prescribed off-diagonal blocks, and

$$\operatorname{rank} P = k(L-1).$$

Proof. The choice of t gives $Q \succ 0$. The Schur complement of Q in P is zero, so $P \succeq 0$ and $\operatorname{rank} P = \operatorname{rank} Q = k(L-1)$. Only diagonal blocks were selected freely: the leading diagonal blocks are those of tI, and the final diagonal block is $B_0^* Q^{-1} B_0$. Thus every off-diagonal block remains exactly (7). □

Positive semidefinite completion by a zero Schur complement is elementary and belongs to the general matrix-completion toolkit [6]. Its consequence here is less obvious: arbitrary boundary subspace data always admit a passive lossless interpolant with a state count linear in both k and L.

Proof of the upper bound in Theorem 1.1. Apply Theorem 3.1. Its matrix obeys (8); the diagonal blocks require no check because $|\zeta_i| = 1$ and $W_i^* W_i = I_k$. Theorem 2.1 produces an interpolant of degree at most $k(L-1)$. $\square$

The construction is algorithmic. It needs one Hermitian eigenvalue bound, one positive linear solve, a factorization of P, and a unitary completion of the lurking isometry. It does not optimize conditioning: increasing t preserves rank but can enlarge stored-energy scales. Degree and numerical conditioning are different resources.

4 The span obstruction and exact generic law

We use two elementary inertia facts. If $P \succeq 0$ has rank q, then

$$\nu_-(P - Z^* P Z) \le q, \tag{10}$$

because the negative inertia of a difference $P - R$ cannot exceed $\operatorname{rank} R$, here $R = Z^* P Z$. Also, adding a positive matrix of rank at most k can remove at most k negative directions. Since $G \succeq 0$,

$$\nu_-(J_L \otimes I_k - G) \ge \nu_-(-G) - k = \operatorname{rank} G - k. \tag{11}$$

Proof of the lower bound in Theorem 1.1. Let $S(z) = D + zC(I - zA)^{-1}B$ be a minimal realization of degree d. The resolvent identity gives

$$S(z) - S(w) = (z - w)C(I - zA)^{-1}(I - wA)^{-1}B.$$

Hence every column of $(S(\zeta_i) - S(\zeta_1))X$ lies in the common space $\operatorname{ran} C$, of dimension at most d. Therefore

$$\mathcal{Y}_1 + \cdots + \mathcal{Y}_L \subseteq \mathcal{Y}_1 + \operatorname{ran} C, \qquad r \le k + d.$$

This proves $d \ge r - k$ directly. Equivalently, the Pick matrix from Theorem 2.1 obeys $\nu_-(P - Z^* P Z) \le \operatorname{rank} P \le d$, while (11) gives $\nu_-(P - Z^* P Z) \ge r - k$. The two arguments audit the same obstruction from realization and interpolation viewpoints. $\square$

Proof of Theorem 1.2. Direct-sum position gives $r = Lk$, so the two sides of (3) coincide. Full column rank of the concatenated frame W is an open condition. Its failure is the zero set of all Lk-minors; when $N \ge Lk$, at least one such minor is nonzero for an orthogonal table, so the full-rank locus is dense. $\square$

The conclusion is stronger than "most targets need many states." It identifies the exact maximum and proves that this maximum occupies an open dense set. Orthogonal targets supply a transparent point in that set, not the mechanism behind the law.

5 Colliding targets and exact-data discontinuity

Let $E_0, E_1, \ldots, E_L$ be mutually orthogonal isometric k-frames in $\mathbb{C}^{(L+1)k}$. For $\epsilon \ge 0$ define

$$W_i(\epsilon) = \frac{E_0 + \epsilon E_i}{\sqrt{1 + \epsilon^2}}, \qquad \mathcal{Y}_i(\epsilon) = \operatorname{ran} W_i(\epsilon). \tag{12}$$

For distinct i, j,

$$W_i(\epsilon)^* W_j(\epsilon) = \frac{1}{1 + \epsilon^2} I_k. \tag{13}$$

Thus all principal angles tend to zero as $\epsilon \downarrow 0$.

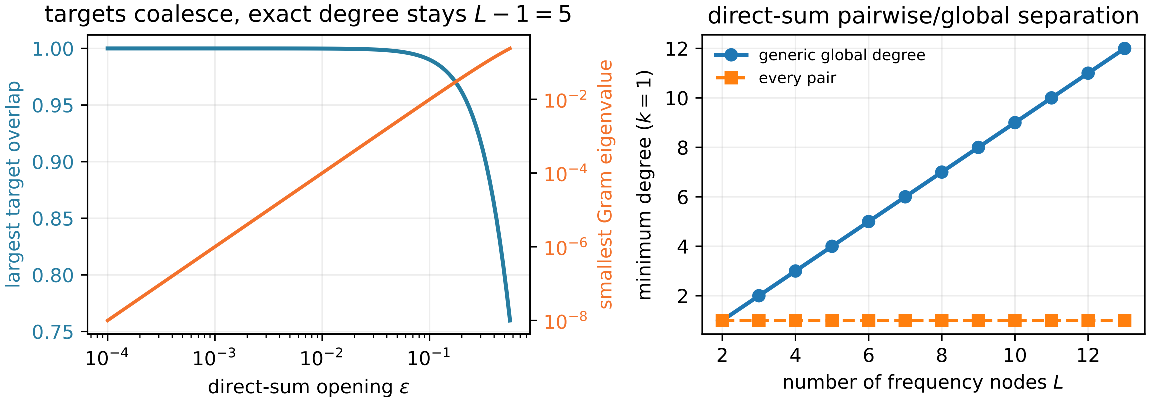

Figure 1: Left: six target lines approach coincidence while the exact minimum degree remains five for every positive opening. The Gram margin closes quadratically. Right: every pair costs one state, whereas the generic global task costs $L-1$. Curves show exact formulas, not fitted asymptotics.

Corollary 5.1 (Maximal memory at arbitrarily small angles)**.** *For the family* (12)*,*

$$d_{\min}(\epsilon) = \begin{cases} k(L-1), & \epsilon > 0, \\ 0, & \epsilon = 0. \end{cases} \tag{14}$$

The block Gram spectrum is

$$\left\{\frac{L+\epsilon^2}{1+\epsilon^2}\right\}^k \cup \left\{\frac{\epsilon^2}{1+\epsilon^2}\right\}^{k(L-1)}. \tag{15}$$

Proof. The private E_i components make the Lk columns of $[W_1(\epsilon)\ \cdots\ W_L(\epsilon)]$ independent for every $\epsilon > 0$. Apply Theorem 1.2. At zero all targets equal $\operatorname{ran} E_0$, and a constant unitary realizes the table. Finally,

$$G_\epsilon = \frac{J_L + \epsilon^2 I_L}{1+\epsilon^2} \otimes I_k,$$

whose spectrum is (15). □

There is no contradiction with continuity of approximate interpolation. The feasible set changes discontinuously only because equality at every node is exact. If a tolerance is admitted, a constant device approximates (12) increasingly well. The small Gram eigenvalue in (15) is precisely the statistical price of distinguishing an expensive positive opening from a free collision.

6 A gauge-invariant finite-error certificate

Let $\Pi_i = W_i W_i^*$ be the orthogonal projector onto $\mathcal{Y}_i$ and set

$$H = \sum_{i=1}^{L} \Pi_i. \tag{16}$$

This observable is independent of all frame gauges, and $\operatorname{rank} H = \dim \sum_i \mathcal{Y}_i$. Suppose Hermitian projector estimates satisfy $\|\widehat{\Pi}_i - \Pi_i\| \le \epsilon_i$. With $\widehat{H} = \sum_i \widehat{\Pi}_i$ and $\eta = \sum_i \epsilon_i$, one has

$$\|\widehat{H} - H\| \le \eta. \tag{17}$$

The error bound may come from deterministic calibration guarantees or a high-probability confidence event. No distributional assumption is needed inside the theorem.

Theorem 6.1 (Noisy span certificate)**.** *Let*

$$q_\eta = \#\{j : \lambda_j(\widehat{H}) > \eta\}.$$

Every exact rational-inner interpolant satisfies

$$\deg_{\mathrm{McM}} S \geq \max\{0, q_\eta - k\}. \tag{18}$$

In particular, if $q_\eta = Lk$, then $d_{\min} = k(L-1)$.

Proof. Weyl's inequality and (17) imply that each eigenvalue of $\widehat{H}$ above η corresponds to a strictly positive eigenvalue of H. Hence $\operatorname{rank} H \geq q_\eta$. Apply the lower half of Theorem 1.1; when $q_\eta = Lk$, also apply its universal upper half. □

The rule is fail-closed: a large η can make it abstain but cannot make it overstate the span on the event (17). For (12), certification requires an error smaller than a quantity of order ϵ^2. Nearly coincident planes are geometrically close but quadratically hard to certify as algebraically independent.

Unlike a block Gram estimate, (16) needs no relative $U(k)$ calibration between frames. A significantly negative eigenvalue $\lambda_{\min}(\widehat{H}) < -\eta$ falsifies the declared projector/error model, because the exact H is positive semidefinite.

7 The exact phase diagram for three target lines

The span lower bound is sharp on the open direct-sum stratum, but not on every singular stratum. The smallest nontrivial example already contains a rigidity phenomenon. Let $L = 3$ and $k = 1$.

Lemma 7.1 (Three-point $\mathbb{CP}^1$ compiler)**.** *Let ℓ_1, ℓ_2, ℓ_3 be distinct lines in a complex two-plane H. There are an orthonormal basis u, v of H, numbers $\alpha, \beta > 0$ with $\alpha^2 + \beta^2 = 1$, and distinct $\tau_1, \tau_2, \tau_3 \in \mathbb{T}$ such that*

$$\ell_i = [\alpha u + \beta \tau_i v]. \tag{19}$$

Either cyclic orientation of the triple (τ_1, τ_2, τ_3) may be chosen. Consequently, for any distinct boundary nodes $\zeta_1, \zeta_2, \zeta_3$, the orientation can be chosen so that a disk automorphism b satisfies $b(\zeta_i) = \tau_i$ for all i.

Proof. Identify $\mathbb{P}(H)$ with the Bloch sphere by sending a unit vector $c_1 e_1 + c_2 e_2$ to

$$\big(2\operatorname{Re}(\overline{c_1} c_2), 2\operatorname{Im}(\overline{c_1} c_2), |c_1|^2 - |c_2|^2\big) \in S^2.$$

Three distinct points of S^2 cannot lie on one Euclidean line, because a line meets the sphere in at most two points. They therefore determine a unique affine plane $n \cdot x = h$. This plane is not tangent, so $|h| < 1$. A rotation in $SO(3)$ sends n to the north pole; lifting it through $SU(2) \to SO(3)$ gives a constant unitary change of basis in H. The three rotated points now have third coordinate h, hence representatives

$$\alpha u + \beta \tau_i v, \qquad \alpha = \sqrt{(1+h)/2}, \quad \beta = \sqrt{(1-h)/2}, \quad |\tau_i| = 1.$$

They have distinct τ_i because the original lines were distinct. Replacing the oriented normal n by $-n$ reverses the induced circle orientation. Choose the orientation matching the cyclic order of the ζ_i. Finally, $\operatorname{Aut}(\mathbb{D})$ acts simply transitively on ordered boundary triples with a fixed cyclic orientation, so a unique degree-one Blaschke factor maps ζ_i to τ_i. □

Theorem 7.2 (Three-line phase diagram)**.** *For three target lines $\mathcal{Y}_1, \mathcal{Y}_2, \mathcal{Y}_3$ at distinct boundary nodes,*

$$d_{\min} = \begin{cases} 0, & \mathcal{Y}_1 = \mathcal{Y}_2 = \mathcal{Y}_3, \\ 1, & \textit{the three lines are pairwise distinct and span a two-plane}, \\ 2, & \textit{otherwise}. \end{cases} \tag{20}$$

Proof. A constant unitary realizes, and is forced by, the first case. A degree-one rational-inner matrix is, up to constant unitary factors, one rank-one Blaschke–Potapov factor [9]

$$\mathrm{diag}(1, b(z), I),$$

where b is a disk automorphism. Applied to a fixed input line, its output is either constant or has the form

$$[u + \rho b(z)v], \tag{21}$$

for orthogonal vectors u, v and $\rho > 0$. The nonconstant orbit is injective on $\mathbb{T}$ and lies in $\mathrm{span}\{u, v\}$. Thus degree one cannot produce a repeated line followed by a different line, nor three lines spanning dimension three.

Conversely, apply Theorem 7.1 to three distinct coplanar target lines. It supplies the latitude form (19) and a degree-one Blaschke factor with $b(\zeta_i) = \tau_i$. Choose constant unitaries $V, U \in U(N)$ so that a unit generator x of the input line obeys $Vx = \alpha e_1 + \beta e_2$, while $Ue_1 = u$ and $Ue_2 = v$. Then

$$S(z) = U\,\mathrm{diag}(1, b(z), I_{N-2})V$$

has McMillan degree one and maps x into ℓ_i at ζ_i.

If the lines span dimension three, Theorem 1.1 gives degree at least two. If exactly two distinct lines occur, injectivity of every nonconstant degree-one orbit gives the same lower bound. The universal upper bound with $L = 3, k = 1$ supplies degree at most two in both cases. □

The second obstruction is genuinely global but not a rank obstruction. Three distinct coplanar lines cost one state; replacing one by a repeated line can raise the cost to two. Algebraic incidence and interpolation order therefore matter on the rank-deficient boundary of the generic stratum. For four or more scalar nodes, degree-one feasibility additionally imposes Möbius cross-ratio compatibility; the three-node case is special because an oriented boundary triple has no remaining modulus.

8 Reproducible synthesis and stress tests

The proof is exact and does not depend on floating-point evidence. The computation checks that the proposed completion and compiler behave as claimed away from the orthogonal case. The producer is `research/generic_fanout_memory_certificate.py`; the independent implementation is `verification/verify_generic_fanout_memory.py`. The verifier imports no producer functions.

The deterministic campaign contains five blocks:

1. 512 random arbitrary-node, nonorthogonal target tables with $3 \le L \le 8$ and $1 \le k \le 3$;
2. 20 near-collision fixtures down to $\epsilon = 10^{-3}$;
3. five unitary colligations, checked for interpolation, unitarity, controllability, observability, and the predicted state dimension;
4. 1,024 noisy projector-sum trials, half direct-sum and half rank deficient;
5. four clustered-node fixtures with gaps from 10^{-2} to 10^{-8}.

Across the 512 completion trials, the worst relative Stein residual was 4.23×10^{-16} and the worst relative violation of positive semidefiniteness was 5.06×10^{-16}. Every completion had rank $k(L-1)$. All 512 direct-sum noisy fixtures were certified and none of the 1,024 trials overcertified the true span bound. These counts and exact floating-point outputs are frozen in a JSON certificate rather than copied from plotting code.

Close nodes are an exact nonissue but a numerical stressor: the entries of the Pick completion scale inversely with the node gap. The clustered-node fixtures retain relative Stein residuals below the declared tolerance down to 10^{-8}, while recording the floating-point rank separately from the

Table 1: Representative colligations synthesized from dense target Gram matrices. Errors are spectral norms.

L	k	states	interpolation	colligation unitary	$\rho(A)$
3	1	2	$1.72\,10^{-14}$	$1.26\,10^{-14}$	0.7881
4	1	3	$2.16\,10^{-14}$	$3.15\,10^{-14}$	0.9610
3	2	4	$3.05\,10^{-13}$	$1.87\,10^{-13}$	0.9918
4	2	6	$3.28\,10^{-14}$	$3.01\,10^{-14}$	0.8677
5	2	8	$2.14\,10^{-14}$	$6.44\,10^{-14}$	0.8584

exact rank guaranteed by the zero Schur complement. We do not claim reliable double-precision colligation synthesis at the smallest gaps.

The release gate performs three distinct checks: hash verification of frozen artifacts, numerical replay in a temporary directory, and an independent formula-level audit. The immutable bundle is linked at `v2.3-direct-sum-fanout-submission`.

9 Scope, relation to prior work, and limitations

Boundary matrix Schur interpolation, minimal rational interpolation, and degree control have a substantial literature [5, 1, 3, 7, 2]. Scalar boundary interpolation also has sharp degree bounds and minimal-degree theory [8, 11]. Subspace interpolation has also been used to design matrix all-pass filters [10, 4]. Positive semidefinite completion and rank questions are classical in their own right [6]. We do not claim any of these general mechanisms as new.

The contribution is the following conjunction for passive spectral routing: the universal sharp sandwich (3); the exact open-dense law (4); a single explicit completion that proves the upper bound for arbitrary subspaces and nodes; the collision discontinuity (14); its finite-error certification threshold; and the complete singular-stratum classification (20). Relative to the orthogonal special case, the result replaces symmetry by a generic geometric condition and identifies what fails when that condition degenerates.

Several boundaries should be kept explicit.

- The task is exact interpolation at finitely many distinct boundary nodes. Approximate, broadband, or derivative constraints define different optimization problems.
- The transfer is finite-dimensional, square, rational, and inner. Loss, gain, infinite delays, and nonrational distributed systems are outside the theorem.
- McMillan degree counts dynamical states. It does not by itself bound dwell time, bandwidth, component sensitivity, or implementation loss.
- The noisy theorem is conditional on valid operator-norm error bounds for the estimated projectors. The simulations illustrate the logic; they are not a substitute for an experimental noise model.
- Outside direct-sum position, $r - k$ need not be exact. The repeated-line case in Theorem 7.2 is the smallest counterexample. A complete stratification for general L and k remains open.

10 Conclusion

Generic spectral fan-out is maximally expensive. The exact state count $k(L-1)$ does not require orthogonal destinations, large angles, or well separated frequencies. It follows from the algebraic independence of the target planes, and it persists at arbitrarily small positive angles. The universal completion proves that no routing table costs more; the span inertia proves that a direct-sum table costs no less.

The collision family separates exact complexity from geometric distance. Its quadratic Gram margin also turns that separation into an operational statement about certification precision. Finally, the three-line phase diagram shows that span is the complete story only on the generic stratum: at singular configurations, Möbius rigidity supplies an additional memory obstruction. Together these results replace one symmetric example by a global law, a sharp boundary phenomenon, and a reproducible compiler.

Data and code availability. Source, producer, independent verifier, frozen certificate, figure, build logs, and artifact manifest are contained in the repository and immutable release linked in Section 8. The numerical evidence is supporting material; all mathematical claims are proved in the text.

AI-assistance disclosure. The author used a generative AI system for literature triage, symbolic and numerical scripting, adversarial proof review, manuscript organization, and language editing. The mathematical claims were rederived as explicit proofs and deterministic certificates. The author retains responsibility for every definition, proof, citation, computation, and submission decision.

References

[1] Athanasios C. Antoulas, Joseph A. Ball, Jeongook Kang, and Jan C. Willems. On the solution of the minimal rational interpolation problem. *Linear Algebra and its Applications*, 137/138:511–573, 1990.

[2] Joseph A. Ball and Vladimir Bolotnikov. The bitangential matrix Nevanlinna–Pick interpolation problem revisited. In *Indefinite Inner Product Spaces, Schur Analysis, and Differential Equations*, volume 263 of *Operator Theory: Advances and Applications*, pages 107–161. Birkhäuser, 2018.

[3] Joseph A. Ball and Jeongook Kim. Stability and McMillan degree for rational matrix interpolants. *Linear Algebra and its Applications*, 196:207–232, 1994.

[4] Agulla Surya Bharath, Devanshu Singh Gaharwar, Kumar Appaiah, and Debasattam Pal. Design of discrete-time matrix all-pass filters using subspace Nevanlinna–Pick interpolation. *Signal Processing*, 204:108839, 2023.

[5] Vladimir Bolotnikov and Harry Dym. *On Boundary Interpolation for Matrix Valued Schur Functions*, volume 181 of *Memoirs of the American Mathematical Society*. American Mathematical Society, 2006. Memoir no. 856.

[6] Jerome Dancis. Positive semidefinite completions of partial hermitian matrices. *Linear Algebra and its Applications*, 175:97–114, 1992.

[7] Andrea Gombani and Gyorgy Michaletzky. On the Nevanlinna–Pick interpolation problem: Analysis of the McMillan degree of the solutions. *Linear Algebra and its Applications*, 425:486–517, 2007.

[8] William B. Jones and Stephan Ruscheweyh. Blaschke product interpolation and its application to the design of digital filters. *Constructive Approximation*, 3(4):405–409, 1987.

[9] V. P. Potapov. The multiplicative structure of J-contractive matrix functions. *Trudy Moskovskogo Matematicheskogo Obshchestva*, 4:125–236, 1955. In Russian; English translation in AMS Translations, Series 2, Vol. 15.

[10] Paolo Rapisarda and Jan C. Willems. The subspace Nevanlinna interpolation problem and the most powerful unfalsified model. *Systems & Control Letters*, 32(5):291–300, 1997.

[11] Gunter Semmler and Elias Wegert. Boundary interpolation with Blaschke products of minimal degree. *Computational Methods and Function Theory*, 6(2):493–511, 2006.