Exact Memory of Finite Spectral Routing Tables: The Direct-Sum Occupancy Law and Its Singular Strata

Lluis Eriksson

11 August 2026

Abstract

We give a single exact state-counting theory for finite passive spectral routing tables. Fix distinct boundary frequencies, one input k-plane, and a word w whose symbol a requests a target k-plane Y_a exactly n_a times. If the used targets are in direct-sum position—they need not be orthogonal and may have arbitrarily small principal angles—then the minimum McMillan degree among square finite rational-inner interpolants is

$$d_{\min} = k\bigl(L - \min_a n_a\bigr).$$

Thus the rarest target fixes generic passive memory, independently of node spacing and word order.

The lower bound uses block-dual detectors: direct-sum geometry supplies a constant compression that is invertible on one target and annihilates every other target. Its determinant has k forced zeros at every wrong node, while a Blaschke–Potapov minor has no more zeros than the network has states. The upper bound is a matrix spectral compiler. Target-weighted node polynomials form a full-rank polynomial column; matrix Fejér–Riesz factorization normalizes it to a rational-inner column of degree at most the lower bound, and a degree-preserving lossless completion closes the network.

The theorem strictly contains orthogonal fan-out, generic all-distinct fan-out, and repeated orthogonal routing as special cases. Beyond the direct-sum locus we prove a detector-rank hierarchy: every constant k-row detector contributes a weighted sum of its rank deficiencies to a universal degree lower bound. For lines this becomes a maximum weighted hyperplane-occupancy bound and exactly resolves the complete three-line phase diagram. We also give a span sandwich, open-dense genericity, fail-closed noisy certification, an exact collision discontinuity, and the optimal integrated Wigner–Smith delay. A producer and independent verifiers audit 512 nonorthogonal scalar tables, 256 block tables, six near-collision openings, and exact singular-incidence fixtures. All source, frozen evidence, and release metadata are linked visibly in the paper.

Keywords: rational inner matrix; McMillan degree; subspace interpolation; Fejér–Riesz factorization; passive memory; spectral routing; Wigner–Smith delay; collision strata.

1 One table, three formerly separate regimes

A finite lossless discrete-time network is represented by a square rational matrix S, analytic on $\mathbb{D}$ and unitary almost everywhere on $\mathbb{T}$. Its McMillan degree is the state dimension of a minimal conservative realization. Let $\zeta_1, \ldots, \zeta_L \in \mathbb{T}$ be distinct, let $X \subset \mathbb{C}^N$ be a k-plane, and let a word $w = w_1 \cdots w_L$ use the targets $Y_a \subset \mathbb{C}^N$. We impose only subspace data,

$$S(\zeta_i)X = Y_{w_i}, \qquad 1 \le i \le L, \tag{1}$$

so bases and endpoint phases remain free. Every competitor is square, finite rational inner, and regular at the interpolation nodes. Put

$$n_a = \#\{i : w_i = a\}, \qquad n_* = \min_{a \in \mathcal{A}(w)} n_a, \qquad m = |\mathcal{A}(w)|.$$

Table 1: Results unified by Theorem 1.2. "New here" is the dense nonorthogonal, repeated-target regime.

Regime	Geometry	Occupancies	Exact degree
Orthogonal fan-out	L orthogonal targets	all $n_a = 1$	$k(L-1)$
Generic fan-out	L direct-sum targets	all $n_a = 1$	$k(L-1)$
Orthogonal word	orthogonal used targets	arbitrary n_a	$k(L-n_*)$
Unified regime (new)	arbitrary direct-sum targets	arbitrary n_a	$k(L-n_*)$

Earlier papers treated fully distinct orthogonal targets, fully distinct direct-sum targets, and repeated orthogonal targets separately [8, 7, 9]. The missing intersection was the natural one: repeated targets with dense, nonorthogonal Gram matrix. That intersection produces the master law.

Definition 1.1 (Direct-sum target table)**.** The used targets are in direct-sum position if

$$\dim \sum_{a\in\mathcal{A}(w)} Y_a = mk. \tag{2}$$

Equivalently, for any isometric frames W_a of Y_a, the block column $W = [W_1 \ \cdots \ W_m]$ has full column rank.

Theorem 1.2 (Direct-sum occupancy law)**.** *Under* (1) *and Definition 1.1,*

$$\boxed{d_{\min}(w; Y_1, \ldots, Y_m) = k(L - n_*).} \tag{3}$$

The value is independent of the locations of the distinct nodes, the principal angles among the targets, and the order of the letters.

Direct-sum position is weaker than orthogonality. The targets may approach one another arbitrarily closely without leaving the locus. Hence Theorem 1.2 removes the most restrictive geometric hypothesis from the routing-word law.

2 Dual minors and the lower bound

Choose an isometric frame X_0 of X and frames W_a of the targets. Since W has full column rank, it has a left inverse D satisfying $DW = I_{mk}$. Write $D_a \in \mathbb{C}^{k\times N}$ for the block row belonging to symbol a. Then

$$D_a W_b = \delta_{ab} I_k. \tag{4}$$

These block-dual detectors replace orthogonal projections.

For an interpolant define

$$g_a(z) = \det(D_a S(z) X_0). \tag{5}$$

At a node labelled a, the two frames of the same target differ by a unitary gauge, so $g_a(\zeta_i) \neq 0$. At a node labelled $b \neq a$, Eq. (4) makes the entire $k \times k$ compression vanish. Each entry has a zero there, and every determinant term is a product of k entries; hence

$$\operatorname{ord}_{\zeta_i} g_a \geq k \quad \text{whenever } w_i \neq a. \tag{6}$$

Lemma 2.1 (Compressed-minor zero budget)**.** *Let S be a square finite rational-inner matrix of McMillan degree d, regular at the boundary points under consideration. For constant maps $A \in \mathbb{C}^{k\times N}$ and $B \in \mathbb{C}^{N\times k}$, if $g(z) = \det(AS(z)B)$ is not identically zero, then the total multiplicity of its boundary zeros is at most d.*

Proof. Factor S into a constant unitary and d rank-one Blaschke–Potapov factors, with multiplicity [11]. An elementary factor is $E(z) = I + (b(z) - 1)P$. On the kth exterior power,

$$\bigwedge^k E(z) = I + (b(z) - 1)\Pi_P, \tag{7}$$

where Π_P selects wedges containing $\operatorname{ran} P$. Thus every exterior-power matrix coefficient is multiaffine in the d scalar Blaschke coordinates. The determinant g is such a coefficient, with $\bigwedge^k A$ and $\bigwedge^k B$ as constant covector and vector. Clearing the d denominators leaves a numerator of degree at most d. A nonzero numerator has at most d zeros with multiplicity. Regularity excludes a hidden pole at an interpolation node. □

Proof of the lower bound in Theorem 1.2. For each used a, Eq. (5) is nonzero at the n_a nodes carrying a and has a zero of order at least k at each of the other $L - n_a$ nodes. By Lemma 2.1, every interpolant obeys

$$\deg_{\mathrm{McM}} S \geq k(L - n_a). \tag{8}$$

Maximizing over a gives $\deg_{\mathrm{McM}} S \geq k(L - n_*)$. □

The proof is gauge invariant. Changing target frames changes the chosen left inverse but preserves Eq. (4); changing X_0 multiplies every minor by a nonzero constant determinant. Orthogonality was merely one way to obtain detectors, not the source of the integer obstruction.

3 A matrix spectral compiler

We now attain the lower bound. For every used symbol set

$$p_a(z) = \prod_{i:w_i \neq a} (z - \zeta_i), \qquad \deg p_a = L - n_a \leq d := L - n_*. \tag{9}$$

Let $P(z)$ be the $mk \times k$ block column with block $p_a(z)I_k$, and define the target-weighted polynomial column

$$F_0(z) = WP(z) = \sum_a p_a(z)W_a. \tag{10}$$

On $\mathbb{T}$, $P(z)$ has rank k: away from the nodes every p_a is nonzero, and at ζ_i exactly the block $p_{w_i}(\zeta_i)I_k$ survives. Since W is injective, $F_0(z)$ has rank k everywhere on $\mathbb{T}$. Therefore

$$R(z) = F_0(z)^* F_0(z) \succ 0, \qquad z \in \mathbb{T}. \tag{11}$$

The matrix Fejér–Riesz theorem gives an outer $k \times k$ polynomial H, of degree at most d, such that

$$R(z) = H(z)^* H(z), \qquad z \in \mathbb{T}, \tag{12}$$

with $\det H$ zero-free in $\mathbb{D}$ [6]. Consequently

$$F(z) = F_0(z)H(z)^{-1} \tag{13}$$

is an $N \times k$ rational-inner column. At $z = \zeta_i$, right multiplication by the invertible matrix $H(\zeta_i)^{-1}$ does not change range, while Eq. (10) has range Y_{w_i}. Thus F realizes every subspace constraint and is regular at the nodes.

Lemma 3.1 (Compiler degree)**.** *The rational-inner column F in Eq.* (13) *has McMillan degree at most kd, and it has a square rational-inner completion of the same degree.*

Proof. Write $H(z) = H_0 + H_1 z + \cdots + H_d z^d$ with H_0 invertible, as follows from outerness. The equation $H(z)u(z) = x(z)$ gives a causal order-d vector recurrence after multiplication by H_0^{-1}. Storing the previous d k-vectors gives an explicit block-companion realization with kd state coordinates; the numerator F_0, whose degree is also at most d, reads from the same delay line and adds no state. Equivalently, the Smith–McMillan pole budget, including the point at infinity in the reciprocal convention, is bounded by kd. Rectangular paraunitary functions admit square lossless completions without increasing realization state dimension [1, Theorem 2.1(ii)]. That reference uses the exterior-disk convention; paraconjugating before and after completion returns the disk-analytic convention used here and preserves degree. □

Completion of Theorem 1.2. Apply Lemma 3.1 with $d = L - n_*$. The completed square network has degree at most $k(L - n_*)$, while Eq. (8) gives the opposite inequality. Equality follows. □

The compiler is finite algebra: form node polynomials, assemble F_0, matrix-spectrally factor $F_0^* F_0$, normalize, and complete losslessly. It does not solve a nonconvex network optimization.

For the orthogonal construction the matrix factor is scalar and the rank amplification can be written $V \otimes I_k$. The load-bearing degree step is then explicit:

$$\boxed{\deg_{\mathrm{McM}}(V \otimes I_k) = k \deg_{\mathrm{McM}} V,} \qquad \det(V \otimes I_k) = \det(V)^k. \tag{14}$$

Indeed, winding of the determinant equals McMillan degree for square finite rational-inner matrices. Formula (14) recovers the scalar compiler as a structured special case of the matrix construction.

4 Consequences: genericity, words, and delay

Corollary 4.1 (Open-dense occupancy law)**.** *Fix a word using m symbols. If $N \geq mk$, the locus in* $\mathrm{Gr}(k, N)^m$ *on which Eq.* (3) *holds is open and dense.*

Proof. For any frames, direct-sum position is full column rank of W. Full rank is open. Failure is the common zero set of its mk-minors; one minor is nonzero on an orthogonal table, so the full-rank locus is dense. □

Thus $k(L - n_*)$ is the generic exact degree for a fixed routing word. It is not an exceptional orthogonal value. In particular:

$$\begin{aligned} \text{one used target}: &\qquad d_{\min} = 0, \\ \text{all } L \text{ targets distinct}: &\qquad d_{\min} = k(L-1), \\ m \text{ equally occupied targets}: &\qquad d_{\min} = kL(1 - 1/m). \end{aligned}$$

The law is invariant under permutations of the word. It therefore separates state cost from transition count: for $k = 1$, `AAAB` has two cyclic switches and degree three, while `ABAB` has four switches and degree two.

For a differentiable unitary boundary value define the Wigner–Smith matrix

$$Q(\theta) = -iS(e^{i\theta})^* \partial_\theta S(e^{i\theta}). \tag{15}$$

The argument principle gives $\int_0^{2\pi} \operatorname{tr} Q \, \mathrm{d}\theta = 2\pi \deg_{\mathrm{McM}} S$ [13, 12, 11]. Hence:

Corollary 4.2 (Exact delay resource)**.** *On the direct-sum locus,*

$$\boxed{\min_S \int_0^{2\pi} \operatorname{tr} Q(\theta) \, \mathrm{d}\theta = 2\pi k(L - n_*).} \tag{16}$$

This is a global winding law, not a pointwise bound on peak delay or a unique dwell-time distribution.

5 Outside the generic stratum

Direct-sum failure does not make routing impossible, but the occupancy statistic alone ceases to be complete. A larger family of constant detectors still gives rigorous information. For $D_{\mathrm{det}} \in \mathbb{C}^{k\times N}$ put

$$\rho_a(D_{\mathrm{det}}) = \mathrm{rank}(D_{\mathrm{det}} W_a),$$

which is independent of the chosen frame of Y_a. Call D_{det} admissible if $\max_a \rho_a(D_{\mathrm{det}}) = k$, and define

$$\Delta(\mathcal{Y}, n) = \max_{D_{\mathrm{det}}\ \text{admissible}} \sum_{a\in\mathcal{A}(w)} n_a\big(k - \rho_a(D_{\mathrm{det}})\big). \tag{17}$$

The maximum exists at the level of the finite rank patterns; at least one admissible detector exists because every individual k-plane has a left inverse. Put $r = \dim \sum_{a\in\mathcal{A}(w)} Y_a$.

Theorem 5.1 (Universal span sandwich)**.** *For arbitrary equal-rank targets at L distinct nodes,*

$$r - k \le d_{\min} \le k(L-1). \tag{18}$$

Proof. For a minimal realization $S(z) = D_0 + zC(I - zA)^{-1}B$ of degree q, the resolvent identity places every column of $(S(\zeta_i) - S(\zeta_1))X$ in $\operatorname{ran} C$. Thus $\sum_i Y_{w_i} \subseteq Y_{w_1} + \operatorname{ran} C$ and $r \le k + q$.

For the upper bound choose target frames at the L nodes. Boundary Nevanlinna–Pick interpolation fixes the off-diagonal blocks of a Hermitian Pick matrix while leaving its diagonal blocks free. Set the first $L-1$ diagonal blocks to a sufficiently large common multiple of I_k and choose the final block to make the Schur complement zero. The completion is positive semidefinite of rank $k(L-1)$; the standard lurking-isometry construction yields a regular square rational-inner interpolant of no larger degree [3, 5]. □

Theorem 5.2 (Detector-rank hierarchy)**.** *For arbitrary equal-rank targets at distinct nodes,*

$$\boxed{d_{\min} \ge \max\{r - k, \Delta(\mathcal{Y}, n)\}.} \tag{19}$$

On the direct-sum locus the second term equals the exact lower bound: $\Delta(\mathcal{Y}, n) = k(L - n_*)$.

Proof. Fix an admissible D_{det} and set $g_{D_{\mathrm{det}}}(z) = \det(D_{\mathrm{det}} S(z) X_0)$. At a node labelled a, the endpoint gauge is invertible, hence $\mathrm{rank}(D_{\mathrm{det}} S(\zeta_i) X_0) = \rho_a(D_{\mathrm{det}})$. If an analytic $k\times k$ matrix has rank q at z_0, constant row and column operations reduce its value to $\mathrm{diag}(I_q, 0)$; every nonzero determinant term must then use at least $k - q$ entries vanishing at z_0. Therefore

$$\mathrm{ord}_{\zeta_i}\, g_{D_{\mathrm{det}}} \ge k - \rho_a(D_{\mathrm{det}}).$$

Admissibility makes $g_{D_{\mathrm{det}}}$ nonzero at at least one interpolation node, so the compressed-minor budget in Lemma 2.1 gives

$$\deg_{\mathrm{McM}} S \ge \sum_a n_a\big(k - \rho_a(D_{\mathrm{det}})\big).$$

Maximize over D_{det} and combine with the realization-span argument in Theorem 5.1. In direct-sum position, each block-dual D_a has rank k on Y_a and rank zero on every Y_b, $b \ne a$. Thus $\Delta \ge \max_a k(L - n_a) = k(L - n_*)$; the universal upper bound in Theorem 1.2 forces equality. □

Corollary 5.3 (Weighted hyperplane occupancy)**.** *For line targets ($k = 1$), let $V = \sum_a Y_a$. Then*

$$\Delta(\mathcal{Y}, n) = \max_{\substack{H \subsetneq V\ \text{hyperplane} \\ \exists a:\, Y_a \not\subset H}} \sum_{a:\, Y_a \subset H} n_a. \tag{20}$$

Thus every proper hyperplane containing a heavily occupied part of the table spends that entire occupancy from the state budget.

Proof. For $k = 1$, an admissible detector is a nonzero functional on V that does not annihilate every target. Its rank deficiency is one exactly on the lines contained in its kernel. Conversely every hyperplane in (20) is the kernel of such a functional. $\square$

The detector hierarchy is exact on the direct-sum locus and on several singular strata, but it is not a complete equality theory. For example, four generic lines in a two-plane have hyperplane bound one while projective interpolation may require more than one state. The smallest complete phase diagram is nevertheless fully controlled by Corollary 5.3.

Lemma 5.4 (Three coplanar lines compile in one state). *Let $\zeta_1, \zeta_2, \zeta_3 \in \mathbb{T}$ be distinct and cyclically ordered. If Y_1, Y_2, Y_3 are pairwise distinct lines in one two-plane, there is a square rational-inner matrix of McMillan degree one satisfying $S(\zeta_i)X = Y_i$ for $i = 1, 2, 3$.*

Proof. Identify the projective lines in the target two-plane with $\mathbb{CP}^1$, hence with the Bloch sphere. The three distinct target points lie on the intersection of the sphere with a unique non-tangent affine plane. A rotation of the Bloch sphere, lifted to a constant unitary in $SU(2)$, sends this circle to a latitude. Reversing its normal if necessary makes the latitude phases $e^{it_1}, e^{it_2}, e^{it_3}$ have the same cyclic order as the nodes. In a suitable orthonormal basis the rotated lines therefore have representatives

$$y_i = \cos\alpha\, e_0 + \sin\alpha\, e^{it_i} e_1, \qquad 0 < \alpha < \frac{\pi}{2}. \tag{21}$$

The boundary action of $\mathrm{Aut}(\mathbb{D})$ is triply transitive on oriented triples, so a degree-one Blaschke factor b satisfies $b(\zeta_i) = e^{it_i}$. With $P = e_1 e_1^*$ and $x = \cos\alpha\, e_0 + \sin\alpha\, e_1$, the elementary Potapov factor

$$E(z) = I + (b(z) - 1)P$$

obeys $E(\zeta_i)x = y_i$. Undo the constant rotation and extend by the identity on the orthogonal complement. Constant unitaries add no state, so the completed square interpolant has degree one. $\square$

Theorem 5.5 (Three-line phase diagram). *For three target lines at three distinct boundary nodes,*

$$d_{\min} = \begin{cases} 0, & \textit{all three lines coincide}, \\ 1, & \textit{the lines are pairwise distinct and span a two-plane}, \\ 2, & \textit{otherwise}. \end{cases} \tag{22}$$

Proof. If all lines coincide, a constant unitary gives degree zero. For three pairwise distinct coplanar lines, Corollary 5.3 gives $d_{\min} \geq 1$ and Lemma 5.4 attains one. If the three lines span dimension three, a hyperplane contains any chosen pair but not the third, so (20) gives $d_{\min} \geq 2$. If exactly two geometric lines occur, the repeated line is itself a hyperplane in their two-plane and has occupancy two, again giving $d_{\min} \geq 2$. The upper half of Theorem 5.1 supplies degree two in both remaining cases. $\square$

6 Collision discontinuity and finite-error certification

Let $E_0, E_1, \ldots, E_m$ be mutually orthogonal isometric k-frames and define, for $\epsilon \geq 0$,

$$W_a(\epsilon) = \frac{E_0 + \epsilon E_a}{\sqrt{1 + \epsilon^2}}, \qquad Y_a(\epsilon) = \operatorname{ran} W_a(\epsilon). \tag{23}$$

For every $\epsilon > 0$ the used targets are direct-sum, although all principal angles tend to zero as $\epsilon \downarrow 0$. At $\epsilon = 0$ they coincide.

Corollary 6.1 (Exact collision jump). *For any word using all m labels,*

$$d_{\min}(\epsilon) = \begin{cases} k(L - n_*), & \epsilon > 0, \\ 0, & \epsilon = 0. \end{cases} \tag{24}$$

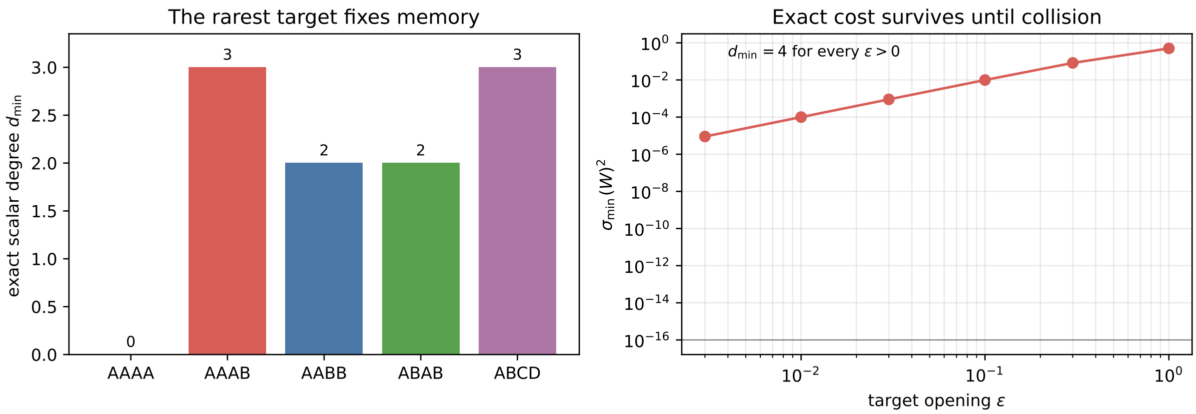

Figure 1: Left: the scalar occupancy law is not ordered by switch count. Right: the direct-sum margin closes continuously in the collision family, while the exact degree stays four for every positive opening and drops only at the singular endpoint. Curves use exact constructions, not fitted asymptotics.

The exact model is therefore discontinuous, even though approximate routing is continuous. This is an algebraic-rank transition, not a claim that a nearly collided table is easy to distinguish experimentally.

There is a gauge-invariant way to certify the correct side. Let Π_a be the projector onto Y_a and $H = \sum_a \Pi_a$. Then $\operatorname{rank} H = \dim \sum_a Y_a$. Suppose estimates obey

$$\|\widehat{H} - H\| \le \eta, \qquad \widehat{H} = \sum_a \widehat{\Pi}_a. \tag{25}$$

Theorem 6.2 (Fail-closed direct-sum certificate)**.** *Let $q_\eta = \#\{j : \lambda_j(\widehat{H}) > \eta\}$. Every exact interpolant satisfies $\deg_{\mathrm{McM}} S \ge \max\{0, q_\eta - k\}$. If $q_\eta = mk$, the targets are certified direct-sum and the exact occupancy law* (3) *follows.*

Proof. Weyl's inequality and Eq. (25) imply $\operatorname{rank} H \ge q_\eta$. Apply the lower half of Theorem 5.1. When $q_\eta = mk$, equality of the span dimension with mk proves direct-sum position and activates Theorem 1.2. □

The test can abstain when uncertainty is large but cannot overstate rank on the declared error event. In Eq. (23), its required precision scales quadratically with the opening, matching the smallest Gram eigenvalue.

7 Executable audit

The numerical layer is deliberately subordinate to the proof. The producer samples dense nonorthogonal target frames, forms a left inverse, checks all dual block isolations, assembles F_0, and verifies strict positivity of $F_0^* F_0$ on a boundary grid. For $k = 1$ it independently computes the outer scalar spectral factor and checks the normalized rational-inner column. Block trials with $k = 2, 3$ audit the matrix precursor and target ranges without pretending that floating-point factorization proves the matrix Fejér–Riesz theorem.

The frozen campaign has 512 scalar and 256 block tables plus six collision openings. The largest dual-isolation residual is 1.17×10^{-11}; the largest target-projector error is 8.72×10^{-11}; the largest wrong dual minor is 1.79×10^{-11}; and the smallest sampled density eigenvalue is 7.82×10^{-9}. In the scalar spectral-factor replay, the largest relative factor residual is 2.37×10^{-8} and the largest boundary isometry residual is 1.78×10^{-5}. The latter is attained in an ill-conditioned fixture and is reported rather than hidden.

Table 2: Theorem-to-artifact map. Paths are relative to the immutable release printed below.

Claim	Executable witness
Dual detectors and spectral compiler	`research/unified_routing_table_certificate.py`
Exact detector-rank incidence fixtures	`research/detector_rank_hierarchy_certificate.py`
Frozen campaign and digest	`results/unified_routing_table_memory/certificate.json`
Independent fail-closed replays	`verification/verify_unified_routing_table_memory.py`, `verification/verify_detector_rank_hierarchy.py`
Collision and occupancy figure	`paper_unified_routing_table_memory/figures/unified_routing_table_memory.pdf`
Paper source and submission record	`paper_unified_routing_table_memory/`

An independent verifier imports no producer code. It recomputes the rarest-occupancy examples and a fresh direct-sum dual fixture, validates the canonical JSON digest, and enforces every numerical tolerance. Both programs fail closed.

8 Positioning, limits, and claims

Minimal rational interpolation, boundary matrix Schur interpolation, and McMillan-degree analysis have an extensive theory [2, 3, 5, 10]. Matrix all-pass design by subspace Nevanlinna–Pick interpolation can also prescribe unitary data and group-delay blocks [4]. Matrix Fejér–Riesz factorization and rectangular lossless completion are classical [6, 1].

The new results are not any one of those ingredients. They are the exact direct-sum occupancy identity (3), including its dual-minor necessity and target-weighted matrix spectral compiler, and the detector-rank hierarchy (19) on arbitrary equal-rank arrangements. A targeted search of the primary literature through 11 August 2026 did not locate the occupancy formula, the statement that least occupancy is a complete generic state statistic, or the weighted rank-deficiency bound in this routing form. This supports priority claims for these precise eliminations, not for rational-inner interpolation, Potapov factorization, or projective geometry in general.

Limitation 8.1 (Exact comparison class). The master theorem assumes distinct nodes, exact subspace data, equal target rank, free endpoint gauges, rationality, and losslessness. Repeated nodes require confluent or jet interpolation. Fixed frames turn phases into data. Unequal ranks create a multiplicity-packing problem. Active or lossy architectures do not obey the inner-function zero and delay budgets.

Limitation 8.2 (Singular geometry). Outside direct-sum position, Eqs. (18) and (19) remain valid but neither occupancy nor the detector-rank statistic alone need determine degree. The paper does not claim a complete equality formula for every intersection lattice; the hierarchy captures all constant determinantal detectors, while compatibility of their upper constructions is a separate problem.

Limitation 8.3 (Conditioning versus degree). Exact McMillan degree is integer-valued and insensitive to node spacing on the theorem's locus. Realization coefficients and spectral-factor conditioning are not. Near node or target collisions, stable numerical synthesis may require higher precision even though the exact degree is unchanged.

The next mathematical frontier is therefore sharper than before: determine when the detector-rank lower hierarchy is attainable and which families of minors can share one state budget on a prescribed intersection lattice. The next physical frontier is a robust approximate law coupling node separation, projector error, and allowed leakage. Neither problem is concealed inside the present exact claims.

9 Conclusion

A finite spectral routing table has a simple generic memory law and a geometric singular lower theory. Direct-sum geometry creates a dual detector for every target; every absence consumes k determinant zeros, so the least visited target asks for the largest budget. More generally, every constant detector charges the occupancy-weighted rank deficiency

$$\sum_a n_a\big(k - \operatorname{rank}(D_{\det}|_{Y_a})\big).$$

A matrix Fejér–Riesz compiler shares exactly the direct-sum budget across all detectors. The equality

$$d_{\min} = k(L - n_*)$$

therefore unifies orthogonal fan-out, generic fan-out, and routing words, while the detector hierarchy, collision family, and constructive three-line lemma expose what survives on singular strata. On the direct-sum locus the same integer is at once a state count, a determinant-winding count, and the optimal integrated Wigner–Smith action divided by 2π.

Immutable reproducibility record. Full release URL (tag included):

`https://github.com/lluiseriksson/finite-sample-spectral-certificates/releases/tag/v2.8-ai-vixra-submission`

The release records the exact target commit, PDF SHA-256, producer commands, independent verifiers, and CI status. Release v2.8 supersedes v2.7 as the submission artifact of this same unified paper. The three predecessor preprints remain citable historical records, but the present article absorbs their scientific scope rather than counting them as additional current results.

References

[1] Daniel Alpay, Palle E. T. Jorgensen, and Izchak Lewkowicz. Characterizations of rectangular (para)-unitary rational functions. *Opuscula Mathematica*, 36(6):695–716, 2016.

[2] 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.

[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] Michael A. Dritschel and James Rovnyak. The operator Fejér–Riesz theorem. *A Glimpse at Hilbert Space Operators*, pages 223–254, 2010.

[7] Lluis Eriksson. Exact memory of direct-sum spectral fan-out. GitHub release, 2026. `https://github.com/lluiseriksson/finite-sample-spectral-certificates/releases/tag/v2.3-direct-sum-fanout-submission`.

[8] Lluis Eriksson. Orthogonal spectral fan-out costs $k(L-1)$ states. GitHub release, 2026. `https://github.com/lluiseriksson/finite-sample-spectral-certificates/releases/tag/v2.1-global-fanout-memory-submission`.

[9] Lluis Eriksson. The rarest target sets passive memory. GitHub release, 2026. `https://github.com/lluiseriksson/finite-sample-spectral-certificates/releases/tag/v2.5-routing-word-memory-final`.

[10] 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.

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

[12] Felix T. Smith. Lifetime matrix in collision theory. *Physical Review*, 118:349–356, 1960.

[13] Eugene P. Wigner. Lower limit for the energy derivative of the scattering phase shift. *Physical Review*, 98:145–147, 1955.