The Exact Four-Level Inverse Commutator Cost

Horn–Littlewood–Richardson facets, rank transitions, and sharp loop synthesis

Lluis Eriksson

10 August 2026

Abstract

For a traceless Hermitian four-level target F, consider the inverse commutator cost

$$\kappa_4(F) = \inf\{\|A\|_{\mathrm{HS}}\,\|B\|_{\mathrm{HS}} : A = A^*,\ B = B^*,\ -i[A, B] = F\}.$$

If $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq \lambda_4$ are the eigenvalues of F, we prove the closed formula

$$\boxed{\kappa_4(F) = \max\{\lambda_1 - \lambda_3,\ \lambda_2 - \lambda_4,\ \lambda_1 - 2\lambda_2 - \lambda_3,\ \lambda_2 + 2\lambda_3 - \lambda_4\}.}$$

The four terms are explicit Horn–Littlewood–Richardson dual certificates. Four spectral chambers attain them through closed self-commutator constructions. On every nonzero spectral stratum, the least rank of an optimal factor is exactly $\max\{n_+(F), n_-(F)\}$, where $n_\pm$ are the positive and negative inertia indices. Thus central interiors use rank two, outer interiors use rank three, and every wall and degeneration is classified as well. More generally, in every finite dimension the inverse cost is one half of a finite Horn linear program for the squared singular values of one matrix. It is consequently a polyhedral, positively homogeneous spectral function and is rational on rational data.

The formula sharpens the universal trace-norm bound:

$$1 \leq \frac{\kappa_4(F)}{\|F\|_1/2} \leq 2.$$

The lower equality holds exactly for sign-paired spectra $(u, v, -v, -u)$, while the upper equality holds only on the two spectral spikes proportional to $(3, -1, -1, -1)$ and $(1, 1, 1, -3)$. Finally, for a balanced three-kick four-level loop with geometric Hamiltonian F, the exact gauge-invariant Hilbert–Schmidt perimeter is

$$\inf S_2^2 = 12\sqrt{3}\,\kappa_4(F).$$

An exact 41-inequality certificate, explicit symbolic constructors, and an independent deterministic linear-program campaign accompany the paper.

1 The inverse question and the novelty boundary

Every trace-zero complex matrix is a commutator, and every finite-dimensional trace-zero Hermitian matrix is a self-commutator [1, 2]. Those existence statements do not answer the inverse metric question: for a *fixed* Hermitian target, what is the least Hilbert–Schmidt size of Hermitian factors that generate it? We study

$$\kappa_d(F) := \inf\{\|A\|_{\mathrm{HS}}\,\|B\|_{\mathrm{HS}} : A = A^*,\ B = B^*,\ -i[A, B] = F\}, \tag{1}$$

where $F = F^* \in M_d(\mathbb{C})$ and $\operatorname{Tr} F = 0$. The norm is unnormalized: $\|X\|_{\mathrm{HS}}^2 = \operatorname{Tr}(X^*X)$. Scalar parts of A and B do not affect the commutator and can only increase the norm, so optimizers may be taken traceless.

The nearby quantitative literature asks different questions. Böttcher–Wenzel inequalities bound a commutator *from above* in terms of its factors and classify direct extremizers [9, 11, 12]. Dimension-dependent commutator theorems bound an operator–Hilbert–Schmidt mixed cost for arbitrary trace-zero matrices [13, 14]. Work on self-commutators establishes existence or forward norm estimates [4, 10, 15]. In contrast, (1) prescribes the whole Hermitian target, requires both factors Hermitian, and minimizes their HS–HS product. To the best of our knowledge, no previous closed spectral formula is available for this four-dimensional inverse optimization problem.

Our engine is classical. Horn's theorem and its representation-theoretic completion characterize spectra of sums of Hermitian matrices [3, 5, 6, 7, 8]. We do not claim that cone or its inequalities. The new step is to turn the prescribed self-commutator problem into a particular Horn linear program, solve its four-level dual exactly, and construct equality matrices on every resulting chamber.

Relation to the ten-paper source chain

This is a new eleventh manuscript, not a replacement for the ten submitted source papers. The directly relevant predecessor [16] solved (1) for $d = 3$ and linked it to balanced three-kick holonomy. The present result is not a formal dimension increment: the qutrit maximum of two adjacent gaps fails in dimension four, two new non-gap facets appear, coherent rank-two constructors replace weighted shifts on central chambers, and the optimal rank undergoes a chamber transition. The other source manuscripts motivate forgotten-order and reflection-positive architectures but are not assumptions of the proofs below.

2 A Horn linear program in every finite dimension

The first reduction replaces two Hermitian factors by one arbitrary matrix.

Lemma 2.1 (Balanced self-commutator reduction)**.** *For every traceless Hermitian $F \in M_d(\mathbb{C})$,*

$$2\kappa_d(F) = \min\{\|C\|_{\mathrm{HS}}^2 : CC^* - C^*C = 2F\}. \tag{2}$$

Proof. Put $C = A + iB$. Direct calculation gives

$$CC^* - C^*C = -2i[A, B], \qquad \|C\|_{\mathrm{HS}}^2 = \|A\|_{\mathrm{HS}}^2 + \|B\|_{\mathrm{HS}}^2 .$$

Rescaling $A \mapsto sA$ and $B \mapsto s^{-1}B$ preserves the commutator and minimizes the displayed sum at

$$2\,\|A\|_{\mathrm{HS}}\,\|B\|_{\mathrm{HS}} .$$

Conversely, the Hermitian real and imaginary parts of every feasible C give feasible factors. Existence and attainment also follow from Theorem 2.2 below. □

We recall the Horn notation only to fix conventions. Let T_r^d be the triples (I, J, K) of r-subsets of $\{1, \ldots, d\}$ with positive Littlewood–Richardson coefficient. For decreasing spectra $\alpha, \beta, \gamma \in \mathbb{R}^d$ with matching trace, there exist Hermitian P, Q such that $\operatorname{spec} P = \alpha$, $\operatorname{spec} Q = \beta$, and $\operatorname{spec}(P + Q) = \gamma$ exactly when

$$\sum_{k\in K} \gamma_k \le \sum_{i\in I} \alpha_i + \sum_{j\in J} \beta_j \quad \text{for every } (I, J, K) \in T_r^d, \quad 1 \le r < d. \tag{3}$$

Theorem 2.2 (General inverse-commutator Horn program)**.** *Let $\lambda_1 \ge \cdots \ge \lambda_d$ be the eigenvalues of a traceless Hermitian F. Define*

$$\alpha(p) = (p_1, \ldots, p_{d-1}, 0), \qquad \beta(p) = (0, -p_{d-1}, \ldots, -p_1), \qquad \gamma = 2\lambda.$$

Then

$$\boxed{\kappa_d(F) = \frac{1}{2} \min_{p_1\ge\cdots\ge p_{d-1}\ge 0} \left\{ \sum_{j=1}^{d-1} p_j : \ (\alpha(p), \beta(p), \gamma) \text{ obey } (3) \right\}}. \tag{4}$$

In particular κ_d is a unitary-invariant, positively homogeneous, continuous polyhedral function of the ordered spectrum. For rational λ, the minimum is rational and has exact rational primal and dual certificates.

Proof. For feasible C, put $P = CC^*$ and $Q = C^*C$. They are positive semidefinite, have the same decreasing spectrum $p_1 \geq \cdots \geq p_d \geq 0$, and satisfy $P - Q = 2F$. If $p_d > 0$, then $P - p_d\mathbf{1}$ and $Q - p_d\mathbf{1}$ remain positive, isospectral, and have the same difference, while their common trace falls by dp_d. Thus a minimum has $p_d = 0$.

Apply Horn's theorem to $P + (-Q) = 2F$. Its three spectra are precisely $\alpha(p)$, $\beta(p)$, and 2λ. Conversely, if the Horn system is feasible, it supplies positive isospectral matrices

$$P = U \operatorname{diag}(p)U^*, \qquad Q = V \operatorname{diag}(p)V^*, \qquad P - Q = 2F.$$

Then $C = U \operatorname{diag}(\sqrt{p})V^*$ obeys $CC^* = P$, $C^*C = Q$, and $\|C\|_{\mathrm{HS}}^2 = \sum p_j$. Lemma 2.1 proves (4). The feasible set is a nonempty closed polyhedron and the objective has compact sublevel sets, so the minimum is attained. The remaining claims are standard consequences of a finite rational linear program on each ordered spectral chamber. □

Remark 2.3. Equation (4) is an exact algorithm, not yet a compact formula: the number of Horn inequalities grows quickly with d. Dimension four is the first case in which the reduction produces the genuinely new chamber geometry solved next.

3 The exact four-level formula

Let $F = F^* \in M_4(\mathbb{C})$ be traceless with ordered eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq \lambda_4$. Write its adjacent gaps as

$$a = \lambda_1 - \lambda_2, \qquad b = \lambda_2 - \lambda_3, \qquad c = \lambda_3 - \lambda_4. \tag{5}$$

Theorem 3.1 (Exact four-level inverse cost)**.** *The inverse Hermitian-commutator cost is*

$$\boxed{\kappa_4(F) = \max\{L_1, L_2, L_3, L_4\}.} \tag{6}$$

Here

$$\begin{aligned} L_1 &= \lambda_1 - \lambda_3 = a + b, \\ L_2 &= \lambda_2 - \lambda_4 = b + c, \\ L_3 &= \lambda_1 - 2\lambda_2 - \lambda_3 = (3a - c)/2, \\ L_4 &= \lambda_2 + 2\lambda_3 - \lambda_4 = (3c - a)/2. \end{aligned}$$

The minimum is attained on all chamber walls and degeneracies.

3.1 Four exact lower certificates

Put $\mu = 2\lambda$, $p = (p_1, p_2, p_3, 0)$, and $s = p_1 + p_2 + p_3$. Table 1 lists Horn inequalities in the notation $(I; J; K)$. Each row is one instance of (3) for $\alpha = (p_1, p_2, p_3, 0)$ and $\beta = (0, -p_3, -p_2, -p_1)$.

Table 1: Horn inequalities used in the four lower certificates.

Certificate	Horn triples	Consequences
L_1	$(1;1;1)$, $(123;124;124)$	$\mu_1 \leq p_1$, $-\mu_3 \leq p_2$
L_2	$(2;1;2)$, $(123;123;123)$	$\mu_2 \leq p_2$, $-\mu_4 \leq p_1$
L_3	$(1;1;1)$, $(12;14;14)$, $(123;134;134)$	$\mu_1 \leq p_1$, $\mu_1 + \mu_4 \leq p_2$, $-\mu_2 \leq p_3$
L_4	$(3;1;3)$, $(23;12;23)$, $(123;123;123)$	$\mu_3 \leq p_3$, $\mu_2 + \mu_3 \leq p_2$, $-\mu_4 \leq p_1$

Summing within each row and using $\sum \mu_j = 0$ gives

$$s \geq \mu_1 - \mu_3, \quad s \geq \mu_2 - \mu_4, \quad s \geq \mu_1 - 2\mu_2 - \mu_3, \quad s \geq \mu_2 + 2\mu_3 - \mu_4.$$

Since $2\kappa_4 = s_{\min}$ by Theorem 2.2, these are exactly the four lower bounds in (6). Notice that L_3, L_4 are not spectral diameters or adjacent gaps; they first appear through three-facet Horn combinations.

3.2 Chambers and attainment

The nonnegative gap octant is partitioned by the four closed chambers in Table 2. On overlaps the displayed values agree.

Table 2: Optimal squared singular spectra $p = \operatorname{spec}(CC^*)$.

Chamber	optimal p	active cost
$c \le a \le c + 2b$	$(2\lambda_1, -2\lambda_3, 0, 0)$	$L_1 = a + b$
$a \ge c + 2b$	$(2\lambda_1, a - c, -2\lambda_2, 0)$	$L_3 = (3a - c)/2$
$a \le c \le a + 2b$	$(-2\lambda_4, 2\lambda_2, 0, 0)$	$L_2 = b + c$
$c \ge a + 2b$	$(-2\lambda_4, c - a, 2\lambda_3, 0)$	$L_4 = (3c - a)/2$

We now construct C explicitly, proving the upper bound without appealing to numerical Horn feasibility.

Upper outer chamber. If $a \ge c + 2b$, equivalently $\lambda_2 \le 0$, set

$$C = \sqrt{2\lambda_1}\, E_{14} + \sqrt{-2(\lambda_2 + \lambda_3)}\, E_{43} + \sqrt{-2\lambda_2}\, E_{32}. \tag{7}$$

All radicands are nonnegative. Direct multiplication yields

$$CC^* - C^*C = 2\operatorname{diag}(\lambda_1, \lambda_2, \lambda_3, \lambda_4),$$

and

$$\|C\|_{\mathrm{HS}}^2 = 2\lambda_1 - 2(\lambda_2 + \lambda_3) - 2\lambda_2 = 2L_3.$$

Upper central chamber. Suppose $c \le a \le c + 2b$. Equivalently there are $A, U, V \ge 0$ such that

$$(\lambda_1, \lambda_2, \lambda_3, \lambda_4) = (A + 2U + V, A, -A - U, -A - U - V). \tag{8}$$

Define

$$x = -\frac{2\lambda_2(\lambda_2 + \lambda_3)}{\lambda_2 - \lambda_4}, \quad y = 2\lambda_1 - x, \quad z = 2\lambda_2 + x, \quad w = -2\lambda_3 - z. \tag{9}$$

For $F \ne 0$, the denominator is positive. In the parametrization (8),

$$x = \frac{2AU}{2A + U + V}, \quad z = 2A + x, \quad w = 2U - x = \frac{2U(A + U + V)}{2A + U + V},$$

and $y = 2(A + 2U + V) - x$, so all four weights are nonnegative. A direct identity gives $xy = zw$. Now put

$$C = \sqrt{x}\, E_{12} + \sqrt{y}\, E_{14} + \sqrt{z}\, E_{23} + \sqrt{w}\, E_{43}. \tag{10}$$

The only possible off-diagonal entry of $CC^* - C^*C$ is $\sqrt{zw} - \sqrt{xy} = 0$. Its diagonal is

$$(x + y, z - x, -z - w, w - y) = 2(\lambda_1, \lambda_2, \lambda_3, \lambda_4),$$

while

$$\|C\|_{\mathrm{HS}}^2 = x + y + z + w = 2(\lambda_1 - \lambda_3) = 2L_1.$$

The zero target uses $C = 0$.

Lower chambers by reflection. Let J be the antidiagonal permutation matrix and define $F^\sharp = -JFJ$. Its ordered gaps exchange a and c. If D is either upper-chamber construction for $F^\sharp$, then

$$C = JD^*J \tag{11}$$

satisfies $CC^* - C^*C = 2F$ with the same HS norm. This proves attainment on the lower central and outer chambers. The four chambers cover $a, b, c \ge 0$, completing the proof of Theorem 3.1.

4 Coherence and exact optimal rank on every stratum

The central constructor is not merely a longer weighted shift. It superposes two incoming paths so that the off-diagonal terms of CC^* and C^*C cancel exactly. For example,

$$F = \operatorname{diag}\left(\frac{13}{4}, \frac{1}{4}, -\frac{3}{4}, -\frac{11}{4}\right) \tag{12}$$

lies in the upper central chamber and has $\kappa_4(F) = 4$. Formula (10) becomes

$$C = \frac{1}{\sqrt{12}} E_{12} + \sqrt{\frac{77}{12}} E_{14} + \sqrt{\frac{7}{12}} E_{23} + \sqrt{\frac{11}{12}} E_{43}. \tag{13}$$

Exactly, $CC^* - C^*C = 2F$ and $\|C\|_{\mathrm{HS}}^2 = 8$.

Proposition 4.1 (Complete minimal optimal-rank law)**.** *Let $F \neq 0$ and let $n_+(F), n_-(F)$ denote its numbers of strictly positive and strictly negative eigenvalues. Then*

$$\boxed{\min\{\operatorname{rank} C : CC^* - C^*C = 2F,\ \|C\|_{\mathrm{HS}}^2 = 2\kappa_4(F)\} = \max\{n_+(F), n_-(F)\}.} \tag{14}$$

Consequently central interiors have optimal rank two and outer interiors rank three; on sign-changing walls the rank drops to two, and on the nonzero two-level ray $(u, 0, 0, -u)$ it drops to one.

Proof. If $\operatorname{rank} C = r$, then both CC^* and C^*C have rank r. Their difference has at most r strictly positive and at most r strictly negative eigenvalues. Hence every feasible C obeys

$$r \geq \max\{n_+(F), n_-(F)\}.$$

For the reverse inequality, inspect the nonzero entries of the optimal squared singular spectra in Table 2. In a central interior there are two, matching inertia $(2, 2)$; the same holds on the common central wall $a = c$ unless both middle eigenvalues vanish. In an outer interior there are three, matching $(1, 3)$ or $(3, 1)$. On an outer/central sign wall one extreme eigenvalue of the three-term spectrum vanishes, leaving rank two and inertia $(1, 2)$ or $(2, 1)$. When both middle eigenvalues vanish, only one squared singular value remains and the inertia is $(1, 1)$. Reflection covers the lower chambers. Thus an optimal constructor attains the inertia lower bound on every stratum. □

Thus inertia forces the rank, while the Horn optimum supplies constructors attaining it without extra singular modes. Figure 1 shows both the resource tax and the four cost chambers on the normalized gap simplex.

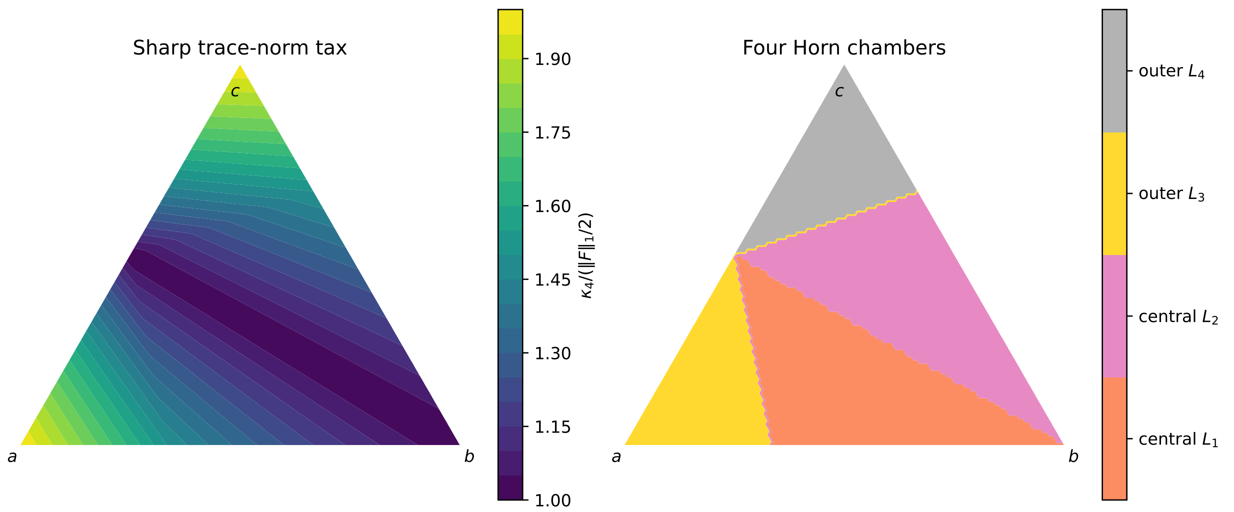

Figure 1: The ordered spectral chamber, represented by nonnegative gaps $a + b + c = 1$. Left: the exact trace-norm tax. Right: the four Horn chambers; central interiors have optimal rank two, outer interiors rank three, and Proposition 4.1 classifies all boundary ranks by inertia.

5 The sharp trace-norm tax

Let

$$P(F) = \frac{\|F\|_1}{2} = \sum_{\lambda_j > 0} \lambda_j$$

for nonzero traceless Hermitian F. The preceding arbitrary-rank theory gives $\kappa_d(F) \geq P(F)$; here we obtain the exact four-level range and equality sets.

Theorem 5.1 (Tax interval and extremizers). *For every nonzero traceless Hermitian $F \in M_4(\mathbb{C})$,*

$$\boxed{1 \leq \frac{\kappa_4(F)}{P(F)} \leq 2.} \tag{15}$$

The lower equality holds exactly when

$$\operatorname{spec} F = (u, v, -v, -u), \qquad u \geq v \geq 0.$$

The upper equality holds exactly when the spectrum is proportional to

$$(3, -1, -1, -1) \quad \textit{or} \quad (1, 1, 1, -3).$$

Proof. It suffices by reflection to use the two upper chambers. In the central upper chamber write $a = c + d$ and $b = (d + e)/2$ with $c, d, e \geq 0$. Direct substitution gives

$$P = c + d + \frac{e}{2}, \qquad \kappa_4 = c + \frac{3d}{2} + \frac{e}{2},$$

so

$$\kappa_4 - P = \frac{d}{2}, \qquad 2P - \kappa_4 = c + \frac{d}{2} + \frac{e}{2}.$$

In the upper outer chamber write $a = c + 2b + t$ with $b, c, t \geq 0$. Then

$$P = c + 2b + \frac{3t}{4}, \qquad \kappa_4 = c + 3b + \frac{3t}{2},$$

and

$$\kappa_4 - P = b + \frac{3t}{4}, \qquad 2P - \kappa_4 = c + b.$$

These four nonnegative differences prove (15); the reflected formulas cover the lower chambers. Equality $\kappa_4 = P$ forces $a = c$, which is equivalent under trace zero to sign pairing. Equality $\kappa_4 = 2P$ forces one of the two outer rays with $b = c = 0$ or $a = b = 0$, giving precisely the stated spikes. ☐

The upper spike is an immediate obstruction to extrapolating the qutrit answer. For

$$F = \operatorname{diag}(3, -1, -1, -1),$$

the spectral diameter is 4 but $\kappa_4(F) = 6$. An extremizer is

$$C = \sqrt{6}E_{14} + 2E_{43} + \sqrt{2}E_{32}, \qquad \|C\|_{\mathrm{HS}}^2 = 12.$$

6 Sharp balanced three-kick synthesis

We now translate the matrix optimum into a finite control resource law. Let h_1, h_2, h_3 be Hermitian kicks in a four-level block, balanced modulo the center. Removing scalar parts gives traceless sides $x_1 + x_2 + x_3 = 0$. For the chosen orientation, the second-order geometric Hamiltonian is

$$H_{\mathrm{geo}} = \frac{i}{2}[x_1, x_2]. \tag{16}$$

The gauge-invariant HS perimeter is

$$S_2 = \sum_{j=1}^{3} \operatorname{dist}_{\mathrm{HS}}(h_j, \mathbb{R}\mathbf{1}) = \sum_{j=1}^{3} \|x_j\|_{\mathrm{HS}}. \tag{17}$$

Theorem 6.1 (Exact four-level triangular action)**.** *Among balanced three-kick loops with* $H_{\mathrm{geo}} = F$*,*

$$\boxed{\inf S_2^2 = 12\sqrt{3}\,\kappa_4(F).} \tag{18}$$

The infimum is attained by an equilateral triangle in the real Hilbert space of traceless Hermitian matrices.

Proof. Let $x_2^{\perp}$ be the HS-orthogonal component of x_2 relative to x_1. The removed component is a scalar multiple of x_1 and hence does not change the commutator. The Euclidean area of the matrix triangle is

$$\Delta = \frac{1}{2}\,\|x_1\|_{\mathrm{HS}}\,\left\|x_2^{\perp}\right\|_{\mathrm{HS}}.$$

Taking $A = x_1$ and $B = -x_2^{\perp}/2$ in (1) and using (16) gives $\kappa_4(F) \le \Delta$. The sharp triangular isoperimetric inequality in a Hilbert space gives

$$\Delta \le \frac{S_2^2}{12\sqrt{3}}.$$

Conversely take factors $-i[A, B] = F$. Orthogonally project B away from A, which preserves the commutator and cannot increase its norm, and rescale $(A, B) \mapsto (sA, s^{-1}B)$ so that they are orthogonal with equal norms. With $\alpha^2 = 4/\sqrt{3}$, define

$$x_1 = \alpha A, \quad x_2 = \alpha\left(-\frac{1}{2}A - \frac{\sqrt{3}}{2}B\right), \quad x_3 = \alpha\left(-\frac{1}{2}A + \frac{\sqrt{3}}{2}B\right).$$

They form a balanced equilateral triangle, have (16) equal to F, and satisfy

$$S_2^2 = 12\sqrt{3}\,\|A\|_{\mathrm{HS}}\,\|B\|_{\mathrm{HS}}.$$

Minimizing over the factors proves (18). The minimum exists by Theorem 2.2. $\square$

The constant in (18) depends on conventions. It uses the squared perimeter $(\sum_j \|x_j\|_2)^2$, not the quadratic energy $\sum_j \|x_j\|_2^2$; the latter has optimum $4\sqrt{3}\kappa_4(F)$. If a full echoed word uses $2H_{\mathrm{geo}}$ as its logged coefficient, the target must be rescaled accordingly. These distinctions prevent a factor-two ambiguity when connecting the theorem to a pulse protocol.

7 Exact and numerical certificates

The mathematical proof above needs only the displayed Horn certificates and direct constructors. The supplementary computation performs a stronger redundant audit:

1. It independently enumerates the Horn triples from Littlewood–Richardson tableaux. For $d = 4$ the counts are 10, 21, and 10 for subset sizes 1, 2, and 3: 41 inequalities in total.
2. It verifies the four dual identities of Table 1 exactly.
3. It parametrizes every closed chamber by three nonnegative variables and proves that all 41 Horn residuals are polynomials with nonnegative rational coefficients.
4. It verifies the outer and coherent central constructors symbolically, including (12)–(13).
5. An independent SciPy/HiGHS campaign solves the full 41-inequality LP for 768 deterministic random spectra. The maximum discrepancy from (6) is 4.45×10^{-16}; the maximum direct constructor residual is 1.08×10^{-15}.

The canonical Horn-triple list has SHA-256

`7bdd2a51639bb83d7cac4d54f4e2874bde8983647f8ff2a08f64787157900241`.

Figure 2 shows the independent LP comparison and the singular-rank transition along a transverse Weyl-chamber slice.

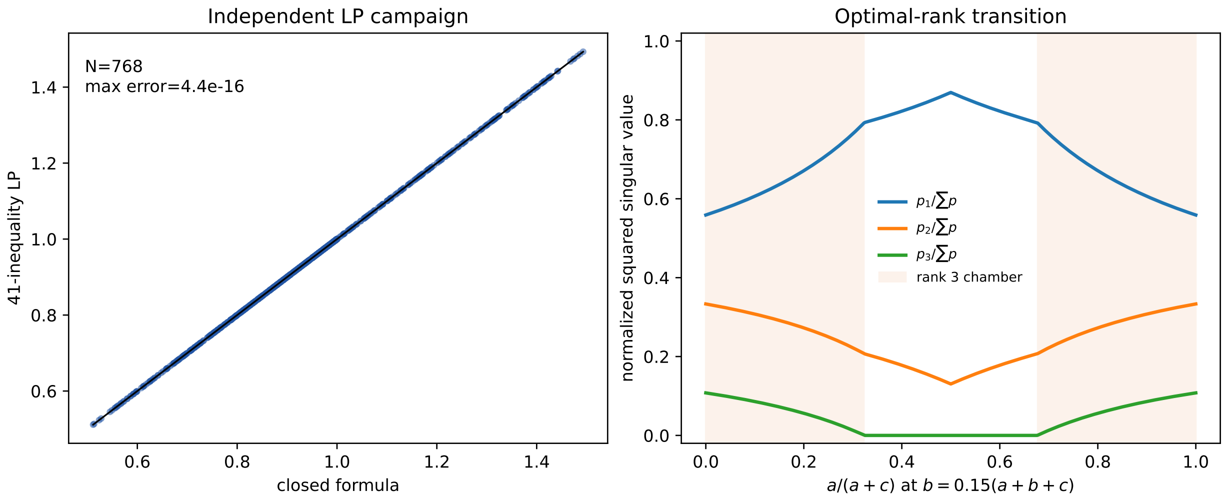

Figure 2: Left: closed formula versus an independent numerical solution of all 41 Horn inequalities. Right: the normalized optimal squared singular values along a slice crossing the two central and two outer chambers; p_3 turns on exactly in the rank-three regions.

The replay command is simply

`python horn_certificate.py  &  python numerical_validation.py`.

The exact certificate uses SymPy; the numerical campaign and figures use NumPy, SciPy, and Matplotlib. No random external data or proprietary solver is required.

8 Scope, limitations, and next frontier

The result is exact but deliberately finite-dimensional. It does not address unbounded canonical commutation relations, locality-constrained factors, operator-norm optimality, infinite factors, or a physical energy model for measurement and pulse duration. The control corollary covers exactly three balanced Hermitian kicks in one four-level block, with the unnormalized Hilbert–Schmidt norm and target convention (16). It is not a controllability theorem under a restricted laboratory gate alphabet.

The general Horn program (4) is finite in every fixed dimension, yet an explicit $d \geq 5$ solution may have many more active facets and higher-rank coherent constructors. Three concrete questions remain:

1. determine the minimal irredundant Horn dual for $d = 5$ and whether a comparably small max formula survives;
2. characterize equality in the trace-norm lower bound in arbitrary dimension directly from the Horn face;
3. impose graph locality or a fixed control algebra and quantify the additional synthesis tax.

The present four-level formula supplies both an exact benchmark and a complete certificate format for those extensions.

Data, code, and AI-assistance statement

All claims are analytic. The accompanying replay bundle contains the LaTeX source, exact Horn/Littlewood–Richardson verifier, deterministic numerical validator, figures, machine-readable results, claim ledger, and SHA-256 manifest. The author directed the research program, selected the problem, reviewed the mathematical claims, and accepts responsibility for the manuscript. OpenAI Codex assisted with repository review, algebraic exploration, literature triage, symbolic and numerical verification, figure generation, and drafting. AI assistance does not replace the explicit proofs or the reproducible certificates supplied here.

References

[1] A. A. Albert and B. Muckenhoupt, "On matrices of trace zero," *Michigan Math. J.* **4** (1957), 1–3. doi:10.1307/mmj/1028990168.

[2] R. C. Thompson, "On matrix commutators," *J. Washington Acad. Sci.* **48** (1958), 306–307.

[3] A. Horn, "Eigenvalues of sums of Hermitian matrices," *Pacific J. Math.* **12** (1962), 225–241. doi:10.2140/pjm.1962.12.225.

[4] P. Fan and C.-K. Fong, "Which operators are the self-commutators of compact operators?," *Proc. Amer. Math. Soc.* **80** (1980), 58–60. doi:10.1090/S0002-9939-1980-0574508-X.

[5] A. A. Klyachko, "Stable bundles, representation theory and Hermitian operators," *Selecta Math. (N.S.)* **4** (1998), 419–445. doi:10.1007/s000290050037.

[6] A. Knutson and T. Tao, "The honeycomb model of $GL_n(\mathbb{C})$ tensor products I: proof of the saturation conjecture," *J. Amer. Math. Soc.* **12** (1999), 1055–1090. doi:10.1090/S0894-0347-99-00299-4.

[7] W. Fulton, "Eigenvalues, invariant factors, highest weights, and Schubert calculus," *Bull. Amer. Math. Soc.* **37** (2000), 209–249. doi:10.1090/S0273-0979-00-00865-X.

[8] A. Knutson, T. Tao, and C. Woodward, "The honeycomb model of $GL_n(\mathbb{C})$ tensor products II: puzzles determine facets of the Littlewood–Richardson cone," *J. Amer. Math. Soc.* **17** (2004), 19–48. doi:10.1090/S0894-0347-03-00441-7.

[9] A. Böttcher and D. Wenzel, "How big can the commutator of two matrices be and how big is it typically?," *Linear Algebra Appl.* **403** (2005), 216–228. doi:10.1016/j.laa.2005.02.012.

[10] G. Nagy, "Abelian self-commutators in finite factors," *Operators and Matrices* **1** (2007), 257–269. doi:10.7153/oam-01-17.

[11] A. Böttcher and D. Wenzel, "The Frobenius norm and the commutator," *Linear Algebra Appl.* **429** (2008), 1864–1885. doi:10.1016/j.laa.2008.05.020.

[12] C.-M. Cheng, S.-W. Vong, and D. Wenzel, "Commutators with maximal Frobenius norm," *Linear Algebra Appl.* **432** (2010), 292–306. doi:10.1016/j.laa.2009.08.008.

[13] W. B. Johnson, N. Ozawa, and G. Schechtman, "A quantitative version of the commutator theorem for zero trace matrices," *Proc. Natl. Acad. Sci. USA* **110** (2013), 19251–19255. doi:10.1073/pnas.1202411109.

[14] O. Angel and G. Schechtman, "The Hilbert–Schmidt version of the commutator theorem for zero trace matrices," *Bull. London Math. Soc.* **47** (2015), 715–719. doi:10.1112/blms/bdv045.

[15] M. Gil', "A sharp bound for the Frobenius norm of self-commutators of matrices," *Linear Multilinear Algebra* **65** (2017), 2333–2339. doi:10.1080/03081087.2016.1273875.

[16] L. Eriksson, "Sharp costs and exact semigroups from forgotten quantum order: qutrit extremality, tetrahedral depolarization, and echo–critical clock laws," submitted manuscript (2026).