

Three-term Weyl operators: central multiplicities, rank laws, and exact list decoding

Lluis Eriksson - 5 September 2026.

Results and scope

Write $X|j\rangle = |j+1 \bmod d\rangle$, $Z|j\rangle = e^{2\pi i j/d}|j\rangle$, and $W_{a,b} = X^a Z^b$. Weyl sparsity s_W counts nonzero coefficients in this single cyclic basis indexed by $\mathbb{Z}_d^2$. A phased Weyl is a Weyl operator multiplied by a scalar of modulus one. All results refer to this fixed cyclic representation; a tensor product of prime-dimensional Pauli bases is a different setting.

For any dimension d , a trinomial $cI + aU + bV$ reduces to central sectors determined by U^q, V^q , where q is the order of their scalar commutator. If u, v are the two labels and $h = |\langle qu, qv \rangle|$, every sector contains $m = d/(qh)$ equivalent cyclic blocks. Each singular sector contributes nullity exactly m . When $abc \neq 0$, there are at most two singular sectors, giving the universal bound $2m$. An example in dimension 36 attains this bound with $m = 3$.

For a noncommuting pair in prime-power dimension, $m = 1$. This yields a complete classification of the ranks of nonzero matrices with at most three Weyl coefficients when $d = p^k$. For $p = 2$, the proper coranks are exactly

$$1, 2, 4, \dots, 2^{k-1}. \quad (\text{A})$$

For odd p , they are exactly

$$\{p^j, 2p^j : 0 \leq j < k\}. \quad (\text{B})$$

The scalar determinant test and the common multiplicity are explicit, including commuting supports and arbitrary composite dimensions. The rank lists (A)-(B) use the additional prime-power arithmetic. No global list of attainable ranks over all composite supports is claimed.

The operational application concerns one use of all d^2 channels $\text{Ad}(W_{a,b})$, pure entangled probes, arbitrary final positive operator-valued measurements (POVMs), and **exact** Schmidt rank. A perfect list decoder reports a set containing the true channel label with probability one. Let $\mathcal{F}_3(d)$ be the reduced density matrices permitting lists of size at most three, and let $\ell_{\min}^{\text{exact}}(d, r)$ minimize the list size over probes of Schmidt rank exactly r . For every $d = 2^k$, $k \geq 3$,

$$\{\rho \in \mathcal{F}_3(d) : \text{rank } \rho < 3d/4\} = \{2P_S/d : S \in \mathcal{H}_d\}, \quad (\text{C})$$

where $\mathcal{H}_d$ consists of exactly thirty half-dimensional subspaces defined below. Thus every rank strictly between $d/2$ and $3d/4$ is forbidden for three labels, and every feasible state below $3d/4$ has a flat half-rank spectrum.

Positive summation changes the answer outside this interval. For every $d = 2^k$, $k \geq 4$, three sparse squares explicitly attain rank $13d/16$. Consequently

$$\ell_{\min}^{\text{exact}}(d, 13d/16) = 3 \quad (k \geq 4), \quad \ell_{\min}^{\text{exact}}(d, 5d/8) = 4 \quad (k \geq 3). \quad (\text{D})$$

In dimension 16, the complete set of ranks admitting at most three labels is therefore

$$\{8, 12, 13, 14, 15, 16\}. \quad (\text{E})$$

Every proof needed for these statements is included below. Section 10 distinguishes the standard representation and factorization tools from their rank and decoding applications here. Earlier internal versions and their separate screening reports are recorded in the companion revision notes.

1. Two elementary ingredients

Character multiplicity. Suppose commuting phased Weyl operators generate a label subgroup H of size h . Every simultaneous character space has dimension d/h ; in particular, $h \mid d$.

To prove this, choose one simultaneous eigenvector and divide each operator by its eigenvalue on that vector. This removes the projective multiplication phases and gives an honest representation $r \mapsto T_r$ of H , with no nonidentity scalar operators. The character projectors are

$$P_\chi = \frac{1}{h} \sum_{r \in H} \overline{\chi(r)} T_r.$$

Every nonidentity Weyl operator has trace zero, so $\text{Tr } P_\chi = d/h$ for every character. All these projectors are nonzero and have the stated rank. This is the standard stabilizer/character averaging argument, reproduced to fix the composite-dimension phase convention.

Circle intersection. If $a, b, c \neq 0$, the equation

$$c + as + bt = 0, \quad |s| = |t| = 1 \quad (1)$$

has at most two ordered pairs (s, t) . Indeed, s must lie on the unit circle and on $|c + as| = |b|$. The second circle has nonzero center $-c/a$, so the two circles are distinct and meet at most twice. Each s determines t .

2. A universal central reduction

Let U, V be phased Weyl operators in the fixed cyclic dimension- d representation, with distinct nonidentity labels u, v . Write $UV = \zeta VU$, let q be the order of ζ , and put

$$H = \langle u, v \rangle, \quad R = \langle qu, qv \rangle, \quad h = |R|. \quad (2)$$

Here $q = 1$ is allowed. For arbitrary d , the equal sector multiplicity is

$$m = \frac{d}{qh} = \frac{dq}{|H|}. \quad (2a)$$

Lemma 1 (central sectors and their multiplicity). The subgroup R is the radical of the commutator pairing on H : its labels commute with every label in H . One has $|H| = q^2 h$. The joint eigenvalue pairs (s, t) of (U^q, V^q) number exactly h , and each joint space $E_{s,t}$ has dimension $d/h = qm$. In particular m in (2a) is a positive integer. On each such space the pair U, V is the orthogonal direct sum of m equivalent q -dimensional cyclic pairs.

Proof. A label $au + bv$ commutes with both u, v if and only if q divides both a, b . This proves the assertion about R . If $au + bv = 0$ as a label, taking its determinant with u and v again shows $q \mid a, b$. Therefore the relation lattice of the map $\mathbb{Z}^2 \rightarrow H$ is contained in $q\mathbb{Z}^2$. The quotient H/qH has order q^2 , and $qH = R$, proving $|H| = q^2 h$.

The operators U^q, V^q are central and have labels generating R . Character multiplicity from Section 1 gives h distinct joint spaces, all of dimension d/h . Their eigenvalue pairs are distinct because the two powers generate this label group after phases are removed.

Fix $E_{s,t}$ and an eigenvalue λ of U on it. The maps V^j , $0 \leq j < q$, send its eigenspace bijectively to the eigenspaces with eigenvalues $\lambda \zeta^j$. These are all the possible eigenvalues because $U^q = sI$. They are distinct, so all have the same dimension m and exhaust $E_{s,t}$. Choose an orthonormal basis $e_1, \dots, e_m$ in the first eigenspace. For each e_l , the vectors $e_l, Ve_l, \dots, V^{q-1}e_l$ span an invariant q -space. The wraparound map is $V^q e_l = te_l$, so these m spaces carry identical cyclic matrices. Orthogonality follows from unitarity and the different U eigenspaces. If $q = 1$, both operators are scalar on the sector and the same statement means m identical one-dimensional pairs. QED.

Theorem 2 (universal trinomial rank rule). Let

$$D = cI + aU + bV, \quad ab \neq 0.$$

No primality assumption on d is needed, and $c = 0$ is allowed in the exact formula. On a central sector $E_{s,t}$,

$$\det(D|_{E_{s,t}}) = [c^q - (-1)^q(a^q s + b^q t)]^m. \quad (3)$$

The sector contributes nullity m if the bracket vanishes and contributes zero otherwise. Consequently

$$\dim \ker D = m |\{(s, t) : a^q s + b^q t = (-1)^q c^q\}|. \quad (4)$$

Equivalently, the rank problem reduces to commuting central powers:

$$\dim \ker D = \frac{1}{q} \dim \ker [c^q I - (-1)^q (a^q U^q + b^q V^q)]. \quad (4a)$$

If all three coefficients are nonzero, then

$$\dim \ker D \leq m \min(2, h) \leq 2m. \quad (4b)$$

Proof. By Lemma 1 it suffices to calculate one cyclic q -dimensional copy. In the U eigenbasis, V is a cyclic weighted shift, and its q nonzero weights have product t . For $q > 1$, the determinant of $cI + aU + bV$ has exactly the diagonal and full-cycle permutation terms. They give

$$\prod_{j=0}^{q-1} (c + a\lambda\zeta^j) + (-1)^{q-1}b^qt = c^q - (-a\lambda)^q + (-1)^{q-1}b^qt.$$

Since $\lambda^q = s$, this is the bracket in (3). A kernel recurrence divides only by b and the nonzero shift weights. Thus a cyclic copy has nullity at most one, including when some diagonal entries vanish or $c = 0$. The determinant condition makes its nullity exactly one when singular. Repeating the copy m times proves (3)-(4). For $q = 1$ the bracket is $c + as + bt$ and each copy is one-dimensional, giving the same conclusion directly.

The central matrix in (4a) is scalar on each sector and has nullity qm there precisely under the same condition. This proves (4a). When $abc \neq 0$, the circle bound of Section 1 applies to $a^qs + b^qt - (-1)^qc^q = 0$. It permits at most two of the h distinct unit-modulus pairs, proving (4b). QED.

The cyclic normal form, scalar central characters and root-of-unity determinant identity are standard Weyl-pair representation machinery. For example, the repeated-eigenvalue cyclic form appears in Farenick, Ojo and Plosker, Lemma 2.2, [primary text](#). The statement here evaluates that machinery for a trinomial in the fixed cyclic Weyl basis, identifies its common multiplicity from the label subgroup, and combines it with the circle count. It is not a claim of a new classification of Weyl-pair representations.

Degenerate central configurations. If either U^q or V^q is scalar, at most one joint pair can solve (4) when $abc \neq 0$, so the bound improves to m . The same improvement holds when R has exponent dividing two. After rephasing, its eigenvalue pairs form a subgroup of the four sign corners. Two distinct zeros differing in one sign force $a = 0$ or $b = 0$; zeros differing in both signs force $c = 0$. This covers $h = 1$, $h = 2$, and the noncyclic order-four center. A cyclic order-four center can have two zeros.

The nonzero constant is essential only for the two-sector bound, not for (3)-(4a). Section 4 exhibits a binomial with six singular sectors. One- and two-term operators are treated separately in the rank classification below.

The equal multiplicity m uses the trace orthogonality of the fixed cyclic Weyl basis. For arbitrary unitaries satisfying $UV = \zeta VU$, central sectors can have different multiplicities. The per-sector cyclic calculation still applies, but the single common value in (2a) is not asserted in that broader setting.

Arithmetic evaluation. If $\delta = u_1v_2 - u_2v_1$ for any integer lifts and $g = \gcd(d, \delta)$, then

$$q = d/g, \quad m = \gcd(g, u_1, u_2, v_1, v_2, \delta/g). \quad (4c)$$

For this optional evaluation, the lattice generated by $u, v, (d, 0), (0, d)$ has index

$$G = \gcd(d^2, du_1, du_2, dv_1, dv_2, \delta), \quad |H| = d^2/G.$$

The index is the standard gcd of the maximal minors of a rank-two integer generating matrix. Substituting in (2a) gives $m = G/g$. Since δ/g is coprime to q , canceling the common q factors gives (4c). The formula is independent of the chosen lifts. The block proof above does not depend on this lattice-index evaluation. In particular, collective primitivity $\gcd(d, u_1, u_2, v_1, v_2) = 1$ suffices for $m = 1$ in any dimension, although it is not necessary.

Corollary 2.1 (prime-power case). If $d = p^k$ and $q > 1$, then $m = 1$ and $h = d/q$. Thus every noncommuting genuine Weyl trinomial in prime-power dimension has nullity at most two. In prime dimension it has at most one, since $h = 1$.

Proof. Write $q = p^\ell > 1$. Then the determinant valuation is $k - \ell < k$, so δ/g is not divisible by p . Formula (4c) gives $m = 1$.

An alternative proof avoids the lattice-index formula. Let S be the minimum coordinate valuation across u, v . After swapping labels if necessary, a determinant-one coordinate change puts $u = (p^S, 0)$ and $v = (c, e)$ with $p^S \mid c, e$. Such a change follows from a Bezout identity for the primitive vector u/p^S ; it is used only for label arithmetic and needs no Clifford implementation. The determinant valuation gives $v_p(e) = k - \ell - s$ and $\ell + 2s \leq k$. The generators qu, qv reduce to the two axes $(p^{\ell+s}, 0)$ and $(0, p^\ell e)$ after subtracting an integer multiple of the first. Their orders multiply to $p^{k-\ell} = d/q$. Hence $h = d/q$ and $m = 1$. The second axis has order one when $s = 0$. QED.

The constant two is attained in dimension eight: take $U = X$, $V = X^2 Z^4$, $c = 1$, $a^2 = 1 - i$, $b^2 = i$. These operators anticommute, and the central condition is $(1 - i)s + is^2 = 1$ for the four eigenvalues of X^2 . Exactly $s = 1, i$ solve it. Section 4 gives a composite example with two singular sectors of multiplicity three, attaining the full bound $2m$.

3. Complete individual rank classification

Theorem 3. Let $d = p^k$, $k \geq 1$. The ranks of nonzero matrices of Weyl sparsity at most three are d , together with $d - n$ for $n \in \mathcal{N}_p$, where

$$\mathcal{N}_2 = \{2^j : 0 \leq j < k\}, \quad \mathcal{N}_p = \{p^j, 2p^j : 0 \leq j < k\} \quad (p \text{ odd}). \quad (5)$$

Proof. Right-multiply by the adjoint of one support Weyl and scale a nonzero coefficient to one. This preserves rank and range, and preserves CC^* up to a positive scalar. A one-term matrix is invertible. A two-term matrix is $I + aU$. If the relative label of U has order m , its eigenvalues are a phase times the m -th roots of unity, each of multiplicity d/m . Thus its nonzero corank, when singular, is d/m , a power of p .

A genuine trinomial is $I + aU + bV$, with distinct nonidentity labels and $ab \neq 0$. If the relative Weyls do not commute, Corollary 2.1 gives corank one or two when singular. For $d = 2$, its one block gives corank at most one. All these possibilities are included in (5).

If the relative Weyls commute, let their label group have order h . The group order is a power of p dividing d , and distinct joint characters have distinct pairs of eigenvalues of U, V . The circle bound gives either one or two zero characters. Their common multiplicity is d/h . Hence the corank is d/h or $2d/h$. When $p = 2$, the presence of three distinct support

labels forces $h \geq 4$; both possibilities are powers of two strictly below d . When p is odd, they have the form in (5).

For attainment of every p^j , use $I - Z^{p^j}$. For odd p , put $m = d/p^j$, $T = Z^{p^j}$, and $\xi = e^{2\pi i/m}$. The polynomial

$$(T - I)(T - \xi I)$$

has exactly three nonzero Weyl coefficients, since $m \geq 3$ is odd and $1 + \xi \neq 0$. Its kernel is the sum of two eigenspaces of T , giving corank $2p^j$. The identity attains full rank. QED.

This is a rank classification, not a complete classification of the kernels or coefficients at each rank. In the dyadic case it says that allowing one individual third Weyl coefficient does not create any new rank beyond those already attained by binomials. That conclusion must not be transferred to a positive sum of several sparse squares; Section 7 gives a counterexample.

4. Composite multiplicities and a necessary coefficient hypothesis

Proposition 4. The factor m in Theorem 2 cannot be discarded. In cyclic dimension 12,

$$C = 2^{1/4}I + X^3 + Z^3 \tag{6}$$

has Weyl sparsity three, noncommuting relative support, rank nine and nullity three.

Proof. Put $U = X^3$, $V = Z^3$. Then $UV = iVU$ and $U^4 = V^4 = I$. Each eigenvalue of U has multiplicity three. Choose a basis of its eigenvalue-one space and apply V, V^2, V^3 ; this decomposes the twelve-dimensional representation into three identical four-dimensional cyclic blocks. On each block, the determinant of $cI + U + V$ is $c^4 - 2$. For $c = 2^{1/4}$ it vanishes, and the invertible-shift recurrence gives a one-dimensional kernel in each block. Hence the total nullity is three. QED.

Here $q = 4$, $h = 1$, and $m = 3$. Thus (6) is a one-sector instance of the universal formula; prime-power noncommuting multiplicity one does not extend to all dimensions. The replay also computes its rank directly over $\mathbb{Q}(i, 2^{1/4})$.

The full two-sector bound is attained with multiplicity greater than one. In dimension 36 take

$$U = X^3, \quad V = iX^{-3}Z^{18}, \quad D = I + U + V. \tag{6a}$$

These Weyls anticommute. Write $S = X^6$, which has order six. Then $U^2 = S$, $V^2 = S^{-1}$, $q = 2$, $h = 6$, and $m = 3$. The sector condition is $s + s^{-1} = 1$. Exactly the two primitive sixth roots of unity solve it, so

$$\dim \ker D = 6 = 2m, \quad \text{rank } D = 30. \tag{6b}$$

Finally, the hypothesis $c \neq 0$ is indispensable for the two-sector bound. In dimension 12, set $U = X$, $V = XZ^6$, and $D = U + V$. Here $q = 2$, $h = 6$, and $m = 1$, while

$$D = X(I + Z^6), \quad \dim \ker D = 6 > 2m. \quad (6c)$$

Indeed Z^6 is an involution with both eigenvalues of multiplicity six. Also $V^2 = -U^2$, so all six central sectors satisfy the exact condition (4), as required. These examples distinguish the universal exact rank formula, its genuine-trinomial bound, and the stronger prime-power specialization.

5. Thirty half-rank squares in every dyadic dimension

For $d = 2^k$, $k \geq 3$, define $\mathcal{H}_d$ as follows:

- Take either eigenspace of each of the three nonidentity order-two Weyls with labels in $\{0, d/2\}^2$.
- For every order-four Weyl with label in $(d/4)\mathbb{Z}_d^2$, take the sum of its eigenvalue- i^j and eigenvalue- i^{j+1} spaces, with j modulo four.

Every subspace has dimension $d/2$. The indicated Weyl representatives have fourth power I . The family contains six subspaces of the first type and twenty-four of the second, after deduplicating inverse order-four generators.

Lemma 5 (low-rank square rigidity). If $C \neq 0$, $s_W(C) \leq 3$, and $\text{rank } C < 3d/4$, then

$$CC^* = cP_S, \quad c > 0, \quad S \in \mathcal{H}_d. \quad (7)$$

Conversely, every such projector is proportional to a sparse square with at most three terms.

Proof. A singular binomial of rank below $3d/4$ must have relative label order two; larger orders give rank at least $3d/4$. Such a binomial is proportional, up to a right unitary, to $I \pm J$ for a phased involution J . Its square is proportional to the corresponding half-rank projector.

For a genuine trinomial, a noncommuting relative support has rank at least $d - 2 \geq 3d/4$, by Corollary 2.1. For commuting support, a group of order at least eight also gives rank at least $3d/4$, by the two-zero-character bound. Thus the group has order four and exactly two characters are zeros.

If this group is $C_2 \times C_2$, absorb phases and write the four values as $1 + a\varepsilon + b\delta$, with $\varepsilon, \delta \in \{\pm 1\}$. Two zeros differing in just one sign force $a = 0$ or $b = 0$; two differing in both signs force the identity coefficient to vanish upon addition. Both are impossible. The group is therefore cyclic of order four.

Let T be its generator, with $T^4 = I$. A right multiplication by a power of T makes the three supported powers I, T, T^2 . The two roots of the resulting quadratic among $\{1, i, -1, -i\}$ cannot be opposite, since their sum would vanish and remove the middle coefficient. They are adjacent. The two remaining roots also form an adjacent pair. After a root rotation and a scalar factor it suffices to consider $p(t) = (t - 1)(t - i)$. Its two nonzero values satisfy

$$|p(-1)|^2 = |p(-i)|^2 = 8.$$

This proves (7). Conversely, $(I \pm J)(I \pm J)^* = 4P_S$ for the involution spaces. For an adjacent order-four pair, take the quadratic whose roots are the complementary pair; its square is $8P_S$. QED.

Lemma 6 (incidence and count). The family $\mathcal{H}_d$ has exactly thirty distinct members. Distinct members intersect in dimension zero or $d/4$. In dimension 8, among the 435 unordered distinct pairs, 207 have zero intersection and 228 have intersection dimension two. For every $d \geq 16$, precisely 15 pairs have zero intersection and 420 have intersection dimension $d/4$.

Proof. An involution projector has Weyl support $\{0, j\}$. An adjacent order-four projector has support $\{0, u, -u\}$, with both nonidentity coefficients nonzero. There are twelve order-four labels, six inverse pairs, and four distinct adjacent projectors per cyclic subgroup. This proves the count and separation from the six involution projectors.

For two different involution labels, the joint group has order four, so fixing both eigenvalues gives intersection dimension $d/4$. For equal labels, the spaces are equal or complementary.

For an involution J and an order-four U , if the label of J equals that of U^2 , an adjacent pair selects one U eigenvalue in each square sector. Its intersection with either J sector has dimension $d/4$. Otherwise $\langle J, U \rangle$ has order eight and two allowed joint characters, again giving $d/4$.

For two order-four generators U, V in dimension $d \geq 16$, all labels under consideration commute: their determinants are multiples of $d^2/16$, hence of d . Their label group has order four, eight or sixteen. In the order-four case, distinct adjacent pairs share either zero or one of the four eigenspaces, each of dimension $d/4$. In the order-eight case, their nonzero squared labels coincide. Rephase V to make $V^2 = U^2$; this only permutes its adjacent-pair family. The eight joint characters satisfy $\lambda^2 = \mu^2$. Each adjacent pair selects one eigenvalue for each square, giving two allowed joint characters and intersection dimension $2d/8 = d/4$. In the order-sixteen case, all sixteen pairs of fourth roots occur, and the two adjacent selections permit four pairs, giving $4d/16 = d/4$. Thus disjointness occurs only for complementary members. Every member has one complement, yielding fifteen unordered disjoint pairs.

In dimension 8, two such order-four Weyls either commute or anticommute. In the commuting case their squared labels coincide, and the order-four/eight argument above applies. In the anticommuting case their squared labels are independent involution labels. The four central character spaces have dimension two. An adjacent-pair projector selects one U eigenline in each such space, and the other projector selects one V eigenline.

Anticommuting invertible operators have no common eigenvector, so these lines differ and the intersection is zero. The six cyclic order-four subgroups form three classes of two according to their squared label. Generators from different classes anticommute. This adds $\binom{3}{2} \cdot 2^2 \cdot 4^2 = 192$ disjoint projector pairs to the fifteen complementary pairs, giving 207. The remaining 228 pairs intersect in dimension two. QED.

6. Operational reduction and the complete low-rank stratum

For completeness, the needed covariance reduction is recalled with its normalization. Represent a pure probe by a full-column-rank matrix $A \in \mathbb{C}^{d \times r}$, with $\rho = AA^*$ and $\text{Tr} \rho = 1$. Vectorization is $|A\rangle\rangle = \sum_{ij} A_{ij} |i\rangle |j\rangle$. A perfect ℓ -list decoder exists exactly when

$$\rho = \sum_j C_j C_j^*, \quad s_W(C_j) \leq \ell. \quad (8)$$

To see necessity, merge effects reporting the same finite list and spectrally refine into vectors $|B_j\rangle\rangle$. Their outcome amplitudes are $\text{Tr}(AB_j^* W_g)$, so each AB_j^* has the reported Weyl support, up to label inversion and phases. POVM normalization gives $\sum_j B_j^* B_j = dI_r$. Setting $C_j = AB_j^* / \sqrt{d}$ proves (8).

Conversely, positivity implies $\text{ran} C_j \subseteq \text{ran} A$. Set $B_j^* = A^+ C_j$. Then $\sum_j B_j^* B_j = I_r$ and $AB_j^* = C_j$. Weyl twirling shows that

$$M_{j,g} = \frac{1}{d} |W_g B_j\rangle\rangle \langle\langle W_g B_j| \quad (9)$$

sum to I_{dr} . For outcome (j, g) , report the translate by g of $\{h : \text{Tr}(C_j W_h) \neq 0\}$. The trace orthogonality of Weyls makes its cardinality $s_W(C_j)$, and it contains every label with nonzero outcome probability. Thus (8) is a physical achievability criterion, not merely a support count.

Theorem 7. Equation (C) holds. Six of the thirty states have minimum fixed-probe list size two; the other twenty-four have minimum fixed-probe list size three.

Proof. Let ρ satisfy (8) with $\ell = 3$ and $r = \text{rank} \rho < 3d/4$. Every summand has range contained in $\text{ran} \rho$, so Lemma 5 applies to every nonzero term. Any two of their half-dimensional ranges, if contained in this same r -space, have intersection dimension at least $d - r > d/4$. Lemma 6 therefore forces them to coincide. Equation (7) and trace normalization give $\rho = 2P_S/d$. Conversely every such state has the sparse-square realization from Lemma 5, and hence a decoder by (9).

For list size two and rank $d/2$, every nonzero sparse summand must itself have rank $d/2$, and its relative label must be an involution. The same range argument forces the state to be one of the first six. All proper-rank probes require at least two labels, because a nonzero one-sparse matrix is invertible. QED.

The cutoff is strict and optimal for this conclusion: rank $3d/4$ already admits two labels, for instance from a binomial whose relative Weyl has order four.

7. Positive summation creates a new rank

Theorem 8. For $d = 2^k$, $k \geq 4$,

$$\ell_{\min}^{\text{exact}}(d, 13d/16) = 3.$$

Construction. Set $m = d/4$, $U = X^m$, $V = Z^m$, and

$$p(t) = (t - 1)(t - i), \quad C_1 = p(U), \quad C_2 = p(V), \quad C_3 = p(UV).$$

These are three-term Weyl matrices. Since $16 \mid d$, U, V commute and generate a label group isomorphic to $C_4 \times C_4$. Its sixteen joint character spaces $E_{x,y}$, where $U = i^x, V = i^y$, have dimension $d/16$.

Put

$$R = \sum_{j=1}^3 C_j C_j^*, \quad \rho = R/(12d). \quad (10)$$

The three kernel conditions are

$$x \in \{0, 1\}, \quad y \in \{0, 1\}, \quad x + y \bmod 4 \in \{0, 1\}.$$

Their simultaneous solutions are exactly

$$(0, 0), (0, 1), (1, 0). \quad (11)$$

Since a positive sum has the intersection of the constituent kernels, $\dim \ker R = 3d/16$, giving rank $13d/16$. Each $C_j C_j^*$ is eight times a half-rank projector, so $\text{Tr } R = 3 \cdot 8(d/2) = 12d$. Thus (10) is normalized and (8) provides a perfect three-label decoder for any purification of ρ .

More explicitly, if $n(x, y)$ counts the failed conditions in (11), then $R|_{E_{x,y}} = 8n(x, y)I$. Among the sixteen cells, the multiplicities for $n = 0, 1, 2, 3$ are 3, 3, 9, 1, respectively. Hence the nonzero spectrum of ρ is

$$\frac{2}{3d} \text{ (3d/16 times), } \frac{4}{3d} \text{ (9d/16 times), } \frac{2}{d} \text{ (d/16 times).} \quad (12)$$

For the lower bound, a proper exact rank r permits two labels if and only if $d - r \mid d$. Here is a proof valid in any cyclic dimension. In a two-sparse factorization (8), positivity forces every nonzero factor to be singular. Normalize a binomial on the right to $\alpha(I - T)$; singularity forces the relative coefficient to have modulus one, so T is a phased Weyl. The common kernel of the positive squares is the intersection of the fixed spaces of these T 's. Since it is nonzero, a common fixed vector forces their scalar commutators to be trivial and forbids nonidentity scalar elements in their generated group. Thus they form an honest commuting representation of a finite label group H . Character averaging gives common-kernel dimension $d/|H|$, proving $d - r \mid d$. Conversely, if $n = d - r$ divides d , the binomial $I - Z^n$ has kernel dimension n ; its normalized square supplies a rank- r state by (8). In the present case $d - r = 3d/16$ does not divide d , so two labels are impossible. QED.

This construction is the decisive counterexample to extending (A) to positive sums. Each individual C_j has rank $d/2$, yet their positive sum has corank $3d/16$, which is not a power of two. It also separates an exact-rank optimization from the rank-budget optimization, which may use a single sparse factor.

Corollary 8.1. In dimension 16, a probe of exact rank r has some perfect decoder with at most three labels if and only if

$$r \in \{8, 12, 13, 14, 15, 16\}.$$

Proof. Theorem 7 excludes all ranks below twelve except eight. Theorem 8 supplies thirteen. The ranks eight, twelve, fourteen and fifteen are attained by the binomials with respective coranks eight, four, two and one. The full-rank maximally entangled probe allows one label. These exhaust the sixteen possible positive ranks. QED.

8. An infinite exact four-label family

Theorem 9. For $d = 2^k$, $k \geq 3$,

$$\ell_{\min}^{\text{exact}}(d, 5d/8) = 4.$$

Proof. The lower bound is Theorem 7, since $d/2 < 5d/8 < 3d/4$. For the upper bound, set $T = Z^{d/8}$, $\omega = e^{2\pi i/8}$, and

$$C = (T - I)(T - \omega I)(T - \omega^2 I).$$

The four coefficients are nonzero. Indeed the two middle elementary symmetric sums are $1 + \omega + \omega^2 = (1 + 1/\sqrt{2})(1 + i)$ and $\omega + \omega^2 + \omega^3 = i(1 + \sqrt{2})$. Exactly three of the eight eigenspaces of T , each of dimension $d/8$, are killed. Thus $\text{rank } C = 5d/8$. Normalize $\rho = CC^*/\text{Tr}(CC^*)$ and use (8). Its exact normalization is $\text{Tr}(CC^*) = 4d(2 + \sqrt{2})$. QED.

The same construction and low-rank obstruction apply uniformly from dimension eight onward.

9. Exact replay, rejected generalizations and limits

Place the two supplied Python files in the same directory and run:

```
python quantum_weyl_verify.py
```

The bounded replay takes about nine seconds after interpreter startup in the recorded environment. It produces `quantum_weyl_certificate.json` beside the script and records the source and companion-checker SHA-256 hashes. It uses integer arithmetic, Gaussian rationals, cyclotomic integer rings, and exact algebraic number fields. No floating-point rank threshold is used. The companion `composite_matrix_checks.py` also runs separately and produces `composite_matrix_certificate.json`.

The certificate checks:

- all 714,096 ordered pairs of distinct nonzero labels in dimensions two through twenty, including commuting pairs, checking $|H| = q^2h$, integrality of m , formula (4c), and the collectively primitive multiplicity-one implication;
- all 1,948,590 noncommuting ordered label pairs in dimensions 2, 3, 4, 5, 7, 8, 9, 16, 25, 27, 32, verifying the prime-power central-group formula by exact lattice-index arithmetic; for dimensions at most eight it separately enumerates the central label group;
- the formal cyclic determinant product for nineteen commutator orders, including even composite orders 6, 10, 12, 18, 20, 24, 36, in integer cyclotomic arithmetic;
- all 54 labelled sparse-square factorizations, all 30 distinct projectors, and all 435 distinct projector pairs in dimensions eight and sixteen;
- the rank $13d/16$ construction by direct exact matrices for $d = 16, 32$, including all sixteen character projectors, their ranks, the complete spectrum and trace normalization;
- the four-term rank $5d/8$ construction for $d = 8, 16, 32, 64$, including the exact nonzero coefficients and zero positions;
- seven direct exact composite-dimension matrix cases covering $q = 1$, one scalar central power, the exponent-two center, two singular sectors, multiplicity three, and the zero-constant negative control. These cases enumerate label groups by closure and compare the rank of the original matrix with the commuting central polynomial in (4a). The resulting $(d, \text{rank } D)$ pairs are (12, 9), (12, 9), (24, 21), (12, 10), (36, 30), (12, 6), (12, 6).

The finite checks do not quantify over arbitrary coefficients or prove statements for untested dimensions. Those quantifiers are supplied by Sections 1-8. The determinant formula avoids division by possibly zero diagonal entries; this is essential in the singular cases.

The composite checker uses a separate computational method and does not import the main verifier. Both were written by the constructing agent; computational separation is not represented as a separate human or agent review. A prior version received separate internal mathematical screening. Its report and the scope of any subsequent screening are recorded outside the article, so that the evidence is tied to the precise source version.

Three broader claims are refuted by the exhibited examples: a multiplicity-free noncommuting nullity bound of two in arbitrary dimension, a two-sector bound when the constant coefficient vanishes, and a power-of-two corank law for all positive sums of three-sparse squares in dyadic dimension. The correct general statement is the central rank formula (4).

No classification of all higher-rank three-list states for arbitrary dyadic dimension is asserted. In particular, (C) controls ranks strictly below $3d/4$; it does not describe the remaining convex set. Nor is a global list of attainable individual ranks for all composite supports derived from the universal central formula. Approximate decoding, mixed probes, and multiple channel uses are outside the scope of the present exact pure-probe results.

10. Antecedents and contribution map

Result or tool	Antecedent inspected	Application in this article
Perfect list decoding and sparse positive factorizations	Johnston, Lovitz, Russo and Sikora, <i>The complexity of perfect quantum state classification</i> , Lemma 9 and Corollary 11, primary full text	Equation (8) gives the Weyl-covariant normalization. No new general factor-width equivalence is claimed.
Equal character multiplicity and stabilizer dimension	Gheorghiu, <i>Standard Form of Qudit Stabilizer Groups</i> , Theorem 1, primary PDF	The trace proof is standard. Lemma 1 evaluates the two-generator central subgroup in the fixed cyclic Weyl representation for arbitrary dimension.
Repeated cyclic blocks for Weyl pairs	Farenick, Ojo and Plosker, <i>Universality of Weyl Unitaries</i> , Lemmas 2.1-2.2, primary PDF	The cyclic normal form is standard. Theorem 2 applies it to rank, with the explicit common multiplicity $d/(qh)$ and the two-circle count; no new representation classification is claimed.
Root-of-unity cyclic representations	Kim, <i>Finite dimensional quantum Teichmüller space from the quantum torus at root of unity</i> , Section 2.1, Propositions 2.5-2.6, primary PDF	Cyclic blocks and central characters are prior machinery. The paper uses odd root order. Our even/odd block signs and multiplicity calculation are proved directly; no even-order theorem is inferred from its odd-order convention.
Quantum-binomial center phenomenon	Fock and Goncharov, <i>Cluster ensembles, quantization and the dilogarithm</i> , Section 3.1.1, primary PDF	Background for the familiar root-of-unity mechanism. The short determinant proof supplies precisely the identity needed here. No claim is made to invent quantum-binomial identities or quantum-torus representations.

Primary-source searches targeted Weyl trinomial rank/nullity, finite Weyl rank uncertainty, finite Gabor three-vector dependence, and root-of-unity quantum-torus representations, including central multiplicities. They located the background machinery above but did not locate the specific prime-power rank list (5), all-dyadic thirty-state stratum, or exact $13d/16$ list value. This bounded negative search does not establish priority. Real-line HRT/Gabor results concern different spaces and quantifiers and were not imported. The companion `quantum_sources.md` records the reading scope and limitations.

The two-label divisibility criterion and the dimension-eight endpoint of the thirty-state and four-label results appeared in earlier internal notes by the same author. Their proofs are included here for self-containment. The contribution developed in this article is the explicit universal rank reduction and its prime-power classification, together with the all-dyadic stratum and the $13d/16$ positive-sum family. Standard representation-theoretic ingredients remain credited to their antecedents.

Lluis Eriksson directed the research programme. OpenAI Codex assisted with conjecture screening, proof development, exact computation, source comparison and drafting. The constructing agent's derivations and exact checks are distinguished from the separate internal review.