Optimal check weight for two-logical-qubit codes of lengths eleven and twelve

Yuxuan Zhang

17 November 2026 (version 1.0)

Abstract

We determine the smallest possible maximum generator weight of binary stabilizer codes encoding two logical qubits at distance at least four, for eleven and twelve physical qubits. In both cases the optimum is five, closing the intervals $[4, 6]$ in the finite-parameter table of Wei et al. The eleven-qubit lower bound follows from exact shadow-enumerator identities, weight-one shortening, and weight parity on an isotropic space. The twelve-qubit lower bound is computer-assisted: every stabilizer is embedded in a self-dual parent from the Danielsen–Parker classification, and an exact enumeration excludes every codimension-two subspace generated by words of weight at most four. This reduction includes degenerate and non-CSS codes and optimizes over all generating bases. Explicit weight-five constructions, proofs of the enumeration criteria, and an executable certificate accompany the result.

1 The two finite optima

A stabilizer group usually admits many generating sets. Even when a code's length and distance are fixed, changing its generators can reduce the largest support of a measured check. The resulting optimization has two levels: one chooses a code and then a generating basis for that code. Wei et al. [1] developed bounds and constructions for this quantity. Their Table 2 leaves the interval $[4, 6]$ at each of $(n, k, d) = (11, 2, 4)$ and $(12, 2, 4)$. We determine both entries.

Work modulo Pauli phases in $V_n = \mathbb{F}_2^{2n}$. An element $v = (x, z)$ has weight $\mathrm{wt}(v) = |\operatorname{supp}(x) \cup \operatorname{supp}(z)|$ and symplectic product

$$\langle (x, z), (x', z') \rangle_{\mathrm{s}} = x \cdot z' + z \cdot x'. \tag{1}$$

A binary $[[n, k]]$ stabilizer is an isotropic subspace $S \subset V_n$ of dimension $n - k$. Consistent stabilizer signs can be assigned to any independent commuting Pauli basis. For $k > 0$, its distance is

$$d(S) = \min\{\mathrm{wt}(e) : e \in S^{\perp} \setminus S\}. \tag{2}$$

Define

$$W(S) = \min_{\mathcal{B} \text{ a basis of } S} \max_{b \in \mathcal{B}} \mathrm{wt}(b), \tag{3}$$

$$W_{\mathrm{opt}}(n, k, d \geq d_0) = \min_{\substack{\dim S = n-k,\ S \text{ isotropic} \\ d(S) \geq d_0}} W(S). \tag{4}$$

Allowing redundant generating sets gives the same minimum: a basis can be selected from any spanning set without increasing its maximum weight.

Theorem 1. *Over all binary stabilizer codes, including arbitrary degeneracy and non-CSS codes,*

$$W_{\mathrm{opt}}(11, 2, d \geq 4) = W_{\mathrm{opt}}(12, 2, d \geq 4) = 5. \tag{5}$$

The lower bound at length twelve uses the published classification of self-dual additive codes through length twelve and the exact finite enumeration specified in Section 4.

The proof uses two different obstructions. At eleven qubits, an enumerator identity limits the number of short stabilizers so severely that a weight-four generating basis would have to consist entirely of even-weight operators. That parity conclusion contradicts the same identity. At twelve qubits, we instead embed each possible stabilizer into a maximal commuting Pauli space. Within each classified parent, the distance condition becomes two elementary tests on a binary two-dimensional subspace. A final rank test decides whether the stabilizer has a weight-four generating basis.

The enumerator framework, parity and shortening arguments, and optimization over generating bases are established tools [1, 2]. The self-dual classification is due to Danielsen and Parker [3]. The contribution here is their application to these two finite optima, together with explicit constructions and complete exclusion certificates. The claim does not concern asymptotic families of low-density parity-check codes.

2 An analytic lower bound at eleven qubits

For an isotropic S let $q(v) = \mathrm{wt}(v) \bmod 2$. At each coordinate one checks

$$q(v + w) = q(v) + q(w) + \langle v, w\rangle_{\mathrm{s}}. \tag{6}$$

Thus $q|_S$ is linear. Nondegeneracy of the symplectic form implies that the affine shadow

$$T = \{h \in V_n : \langle h, s\rangle_{\mathrm{s}} = q(s) \text{ for every } s \in S\} \tag{7}$$

is nonempty. Write A_i, B_i, T_i for the numbers of weight-i elements of $S, S^{\perp}, T$, respectively. All are nonnegative integers, $A_0 = 1$, and $\sum_i A_i = |S|$. Moreover,

$$d(S) \geq 4 \quad \Longrightarrow \quad B_j = A_j \quad (j = 1, 2, 3). \tag{8}$$

This formulation retains low-weight stabilizers and therefore includes degenerate codes.

For completeness, define the quaternary Krawtchouk coefficients by

$$K_j(i) = [z^j](1 + 3z)^{n-i}(1 - z)^i. \tag{9}$$

Summing symplectic characters over a weight sphere gives

$$B_j = \frac{1}{|S|}\sum_i K_j(i)A_i, \qquad T_j = \frac{1}{|S|}\sum_i (-1)^i K_j(i)A_i. \tag{10}$$

Indeed, the indicator of $S^{\perp}$ is the average of $(-1)^{\langle e,s\rangle_{\mathrm{s}}}$ over $s \in S$; the indicator of T has the additional factor $(-1)^{q(s)}$. A zero coordinate of s contributes $1 + 3z$, and a nonzero coordinate contributes $1 - z$. These identities do not require linearity over $\mathbb{F}_4$.

Put $L_5 = A_1 + \cdots + A_5$. Substitution in (10) at $n = 11$, $|S| = 512$, yields the coefficientwise identity

$$\begin{aligned} -L_5 = 10A_0 - \frac{19}{512}\sum_i A_i - 9(B_1 - A_1) - 3(B_2 - A_2) - (B_3 - A_3) \\ + 14T_0 + 4T_1 + 2T_2 + 2A_3. \end{aligned} \tag{11}$$

Using (8) and normalization gives

$$9 - L_5 = 14T_0 + 4T_1 + 2T_2 + 2A_3 \geq 0. \tag{12}$$

We will also use the following inherited ten-qubit obstruction. Its short certificate is included to make the eleven-qubit argument self-contained.

Lemma 2. *No binary* $[[10, 2, d \geq 4]]$ *stabilizer has a basis of weight at most four.*

Proof. For $n = 10$ and $|S| = 256$, the transforms give

$$\begin{aligned} -L_5 = 20A_0 - \frac{7}{64}\sum_i A_i - 12(B_1 - A_1) - 4(B_2 - A_2) - (B_3 - A_3) \\ + 8T_0 + 4T_1 + A_1 + 3A_2 + 3A_3 + A_4. \end{aligned} \tag{13}$$

Consequently

$$8 - L_5 = 8T_0 + 4T_1 + A_1 + 3A_2 + 3A_3 + A_4 \geq 0. \tag{14}$$

Eight independent nonzero generators of weight at most four force $L_5 \geq 8$. Equality then forces $A_1 = A_2 = A_3 = A_4 = 0$, which rules out those generators. □

Proposition 3. *Every binary* $[[11, 2, d \geq 4]]$ *stabilizer satisfies* $W(S) \geq 5$.

Proof. Suppose that nine generators all have weight at most four. They are distinct nonzero stabilizer elements, so $L_5 \geq 9$. Equation (12) forces equality and, in particular, $T_0 = A_3 = 0$.

We first show that $A_1 = 0$. If a weight-one stabilizer exists, a local Clifford operation takes it to Z on one coordinate. Commutation forces all stabilizers to use only I or Z there. Delete that coordinate. The restriction map on S has one-dimensional kernel, so its image has rank eight and still encodes two logical qubits. The images of the alleged generators have weight at most four and span this image.

The shortened code still has distance at least four. A short logical Pauli on it, extended by identity on the deleted qubit, commutes with the original stabilizer. If the extended Pauli lay in the original stabilizer, its restriction would lie in the shortened one, a contradiction. Equivalently, any choice of Z on the deleted coordinate can be removed by multiplying the weight-one stabilizer. The shortened code therefore contradicts Lemma 2.

We have $A_1 = A_3 = 0$. Every generator of weight at most four must now have weight two or four. Linearity of q on S implies $q|_S = 0$. Hence the zero vector belongs to T, giving $T_0 = 1$, contrary to $T_0 = 0$. □

3 Weight-five constructions

Table 1 gives commuting independent Pauli generators for both lengths. The eleven-qubit generators all have weight five. The first three twelve-qubit generators have weight four and the remaining seven have weight five. Direct exact enumeration gives

$$(A_i)_{i=0}^{11} = (1, 0, 0, 0, 0, 9, 42, 105, 135, 115, 78, 27), \tag{15}$$

$$(A_i)_{i=0}^{12} = (1, 0, 0, 0, 3, 12, 20, 108, 243, 260, 204, 132, 41). \tag{16}$$

Their full symplectic duals have $8,192$ and $16,384$ elements. Neither has a logical Pauli of weight below four; they have respectively 60 and 54 logical Paulis of weight four. Thus both distances

[[11, 2, 4]]	[[12, 2, 4]]
XIZIXIXIIZI	IYIIZIIIIXXI
ZYIIZYXIIII	XIIZIIIZZIII
IIXYYZIXIII	IIYIIXYIIIIY
ZIIZZIIIXIZ	YIYZIIIXIXII
ZIXYIIIIYZI	IIZYIZIXIIXI
IZIYIIYIZXI	IYIXIIYZXIII
IIZIIZZYIYI	IIIXZIZYIIIX
IIYIIYIXIXZ	IIIZXZZIIIZI
IXXIYIXIIIY	IXIIIYIIZIYX
	ZIIIIXIIXZYI

Table 1: Explicit generating bases of maximum weight five. Each character acts on one physical qubit.

are exactly four. The companion checker retains Pauli phases and verifies that neither commuting group contains minus identity.

The length-twelve upper bound can also be obtained by appending a fixed qubit to the eleven-qubit code. Appending I to the nine generators and adjoining $I^{\otimes 11}Z$ preserves $k = 2$, distance four, and maximum generator weight five. This degenerate example is included as a positive control. The separate twelve-qubit witness in Table 1 has no stabilizer of weight below four.

4 Excluding weight four at twelve qubits

4.1 Reduction to classified self-dual parents

Any rank-$(n-2)$ isotropic subspace $S \subset V_n$ extends to a Lagrangian subspace C of rank n. Such spaces correspond to self-dual additive codes over $\mathbb{F}_4$. Danielsen and Parker classify all of them through length twelve, up to coordinate permutations and local Clifford operations [3]. Their equivalences preserve both Pauli weight and commutation, hence preserve distance and the optimization in (3).

Every class admits a graph representative

$$C_\Gamma = \{c(t) = (t, \Gamma t) : t \in \mathbb{F}_2^n\}, \tag{17}$$

with Γ symmetric and zero on the diagonal. At $n = 12$, the complete classification contains $1,323,363$ representatives, including decomposable codes. The $1,274,068$ connected representatives alone would not suffice.

Within a fixed parent, every codimension-two subspace is uniquely

$$S_U = \{c(t) : t \in U^\perp\}, \qquad U \le \mathbb{F}_2^n, \quad \dim U = 2, \tag{18}$$

where the perpendicular on the coefficient space uses the ordinary binary dot product. For a physical Pauli $e = (x, z)$, put

$$\sigma(e) = z + \Gamma x. \tag{19}$$

Then $\langle e, c(t)\rangle_{\mathrm{s}} = \sigma(e) \cdot t$, so

$$e \in S_U^\perp \quad \Longleftrightarrow \quad \sigma(e) \in U. \tag{20}$$

Lemma 4. *The subspace S_U has distance at least four if and only if both of the following conditions hold:*

1. *No nonzero element of U is the syndrome of a Pauli of weight one, two, or three.*
2. *Every t with $1 \le \mathrm{wt}(c(t)) \le 3$ is orthogonal to U.*

Proof. For an error of nonzero syndrome, membership in C_Γ is impossible. Such an error is logical precisely when its syndrome lies in U, by (20). For an error of zero syndrome, $z = \Gamma x$ and $e = c(x)$. It lies in S_U exactly when $x \perp U$. These two cases exhaust all errors of weight at most three. □

The second condition allows short stabilizers but excludes short logical operators inside the parent. It is the part of the reduction that handles degeneracy.

For a fixed parent define $H \subset \mathbb{F}_2^n \setminus \{0\}$ to consist of the vectors that are absent from the nonzero syndromes of errors of weight at most three and are orthogonal to all t with $1 \le \mathrm{wt}(c(t)) \le 3$. A two-plane $U = \mathrm{span}\{u, v\}$ passes Lemma 4 exactly when $u, v, u + v \in H$. Choose the usual integer ordering of binary vectors. The test $u < v < u + v$ visits each such plane once.

Lemma 5. *For any cap w, $W(S_U) \le w$ if and only if*

$$\mathrm{rank}_{\mathbb{F}_2}\{t : 1 \le \mathrm{wt}(c(t)) \le w,\ t \perp U\} = n - 2. \tag{21}$$

Proof. The map $t \mapsto c(t)$ is a linear injection. The displayed vectors correspond to every nonzero element of S_U of weight at most w. They contain a basis exactly when they span the full rank $n - 2$. □

It is safe to discard a parent if all of its words of weight at most four span rank less than $n - 2$: restricting to S_U cannot increase that rank. One may also discard a parent whose words of weight at most three span rank greater than $n - 2$, since Lemma 4 forces all these words into S_U. These tests exclude parents by necessary conditions; neither imposes a code ansatz.

4.2 Exact enumeration and its evidence

The source archive of the classification is supplied by Danielsen [4]. The copy used here has SHA-256

`75af3203e86e6ac5726795c84bdacd32f1d13d7c0f1f84b8e3a664724b8d5855.`

The archive contains $45,144$ length-eleven and $1,323,363$ length-twelve representatives. The latter have minimum-distance histogram

$$(N_1, N_2, N_3, N_4, N_5, N_6) = (45144, 615180, 467519, 195456, 63, 1). \tag{22}$$

The archived automorphism orders also reproduce the orbit mass

$$\sum_C \frac{6^n n!}{|\,\mathrm{Aut}(C)|} = \prod_{j=1}^{n} (2^j + 1). \tag{23}$$

These integrity checks agree with the published classification. They do not replace its completeness theorem.

Three exact implementations of the reduction were run. The first enumerates short graph-code words and physical-error syndromes directly. The second uses Gray-code word generation, a Cayley-graph computation of the syndrome ball, and Walsh transforms to count low-weight words in $U^\perp$ before elimination. The companion package contains a compact third implementation written

during review. It extracts graph6 records from the raw archive, decodes their adjacency matrices separately, and computes the syndrome ball as sums of up to three one-site syndromes. Repeated sites are harmless: their Pauli products have support no larger than three. Conversely, every error of support at most three appears in this enumeration.

The third implementation omits the optional short-word-rank pruning test. It still reaches the same $320,777,184$ admissible planes in parents that survive the cap-four rank test. Its complete result is shown in Table 2. No plane has the rank ten required by (21). Among planes with at least ten short words, the maximum rank is only eight: counting short words is not enough to certify a basis.

Quantity	Count or value
Parent representatives processed	$1,323,363$
Parents excluded by cap-four span rank	$734,891$
Admissible planes in the remaining parents	$320,777,184$
Maximum cap-four span rank	9
Planes containing at least ten cap-four words	$564,649$
Maximum rank among those planes	8
Required stabilizer rank	10

Table 2: Complete length-twelve exclusion in the packaged implementation. Both earlier implementations give the same admissible-plane count.

The input count and end-of-file condition are asserted. A known feasible eleven-qubit, cap-five instance is accepted by the same census program. Separate small-graph comparisons check the reduction against direct enumeration of physical errors, including a control that detects omission of the zero-syndrome condition.

Every putative $[[12, 2, d \geq 4]]$ cap-four stabilizer has a self-dual extension, an equivalent classified graph parent, and a plane covered by this enumeration. The negative census therefore excludes every such stabilizer. Together with Proposition 3 and Table 1, this proves Theorem 1.

5 Reproduction and comparison

The source package includes the two constructions, exact enumerator checks, the classification census, and the raw classification input with its provenance. Running `python3 reproduce.py` in an extracted directory checks the packaged hashes, verifies both witnesses and the analytic identities, builds the census with a C++17 compiler, runs a feasible control, and reruns the full twelve-qubit exclusion. The computation uses only exact integer and binary operations. The classification remains a published external dependency; it is not regenerated from first principles.

As of 17 November 2026, the two relevant entries of Wei et al.'s v2 table remain [4, 6]. We also checked the two twelve-qubit, two-logical-qubit, distance-four codes in the public QEC Challenge repository [5]. For each code the span of all words of weight at most four or five has rank six, while the weight-six span has rank ten. Their optimal generating-basis weight is therefore six, not merely six for the displayed generators. The present witnesses and converses close the two finite table intervals.

Research process. The author directed the research and revision. AI assistance was used for candidate derivations, source comparison, code, computational checks, and drafting. The proof

audit and the differently implemented checks reported here were sequential internal work. They do not constitute independent expert review or formal proof-assistant verification.

References

[1] F. Wei, Z. Han, A. Y. He, Z. Li, and Z.-W. Liu, *Theory of low-weight quantum codes*, arXiv:2601.19848v2 (2026), Table 2, Section V, and Appendices H and L.

[2] E. M. Rains, *Quantum shadow enumerators*, IEEE Transactions on Information Theory **45**, 2361–2366 (1999), doi:10.1109/18.796376.

[3] L. E. Danielsen and M. G. Parker, *On the classification of all self-dual additive codes over GF(4) of length up to 12*, Journal of Combinatorial Theory, Series A **113**, 1351–1367 (2006), arXiv:math/0504522v4.

[4] L. E. Danielsen, *Self-dual additive codes over GF(4)*, classification archive, `https://www.codetables.de/larsed/vncorbits/`. The precise input hash is given in Section 4.

[5] Unitary Foundation, *QEC Challenge code repository*, public code data. Files `codes/12-2-4.json` and `codes/12-2-4-b.json`, checked 17 November 2026; immutable Git blob identifiers are included in the source package.