EXACT MINIMAL LENGTHS FOR CONSTANT-EXCITATION QUANTUM CODES: EIGHT QUBITS FOR STABILIZER CODES, TWELVE AND FOURTEEN FOR CSS

KOREA SUPERINTELLIGENCE LABS

Abstract. A quantum code is *constant excitation* (CE) if every computational basis ket occurring in its code space has the same Hamming weight; such a code is immune to collective coherent Z-rotations. We determine the exact minimal length in three cases. First, no $[[n, 1, d]]$ CE CSS code with $d \geq 3$ exists for $n \leq 11$; since Lai, Liou and Ouyang exhibit a $[[12, 1, 3]]$ one, twelve is the exact minimum for a single logical qubit, and the case $n = 10$ is an existence question their paper states as open. Second, no such code with $k \geq 2$ logical qubits exists for $n \leq 13$, so their $[[14, 3, 3]]$ code is likewise smallest, and fourteen is the exact minimum for more than one logical qubit. Third, dropping the CSS restriction: no $[[n, k, d]]$ constant-excitation *stabilizer* code with $k \geq 1$ and $d \geq 3$ exists for $n \leq 7$, so with the $[[8, 1, 3]]$ code of Plenio, Vedral and Knight — verified here from its printed equations — eight is the exact minimum in the wider class. The last supplies a proof of a minimality that Plenio, Vedral and Knight asserted in 1996 on the strength of an undescribed computer search, after deleting an unsound counting argument between the two versions of their preprint; a literature check found no proof of it beyond its CSS sub-case, which is a lemma of Lai, Liou and Ouyang. Each proof reduces the search space by structural theorems — constant excitation is exactly a containment in a cone V_y, and the global shift may be sorted and complemented — and then runs an exhaustive search over the reduced space. The reductions and the completeness of the enumerations are theorems, not assumptions. We also record, conservatively, an erratum to a published catalogue entry for the eight-qubit code. Every theorem below is machine-checked in Lean 4, the passages from the quantum-mechanical definitions to the finite ones excepted; those are marked where they occur.

1. Introduction

Let $\theta \in \mathbb{R}$ and consider the collective rotation $\exp(-i\theta \sum_{j=1}^{n} Z_j)$ — the error an n-qubit register picks up when every qubit is dephased by the same drifting clock. A superposition of computational basis kets all of Hamming weight w is an eigenvector of this operator for every θ, so it acquires only a global phase. A quantum code whose whole code space consists of such states is therefore immune to collective coherent noise. Following [10] we call an n-qubit state with this property a *w-CE state* and a code whose code space consists entirely of w-CE states a *w-CE code*. CE codes go back to Plenio, Vedral and Knight [12], whose eight-qubit code corrects one general error and has all sixteen of the computational kets occurring in its two logical basis states of Hamming weight 4, so it is 4-CE. Such codes have been studied since as the natural carrier for dual-rail concatenation [11] and for balanced weight-2 Z-stabilizers [7].

Two minimal-length questions have been asked about them, at two levels of generality, and this paper answers both.

The CSS question. Lai, Liou and Ouyang [10, Appendix A.1] generalise the CSS construction to produce CE codes by adding a *global shift*. Given linear codes $C_2 \subset C_1 \subseteq \mathbb{F}_2^n$ and a vector $y \in \mathbb{F}_2^n$, the logical basis states are

$$|j\rangle_L \;=\; |C_2|^{-1/2} \sum_{c \in C_2} \big|\, y + x^{(j)} + c \,\big\rangle, \tag{1}$$

where $x^{(j)}$ runs over coset representatives of C_2 in C_1; the number of logical qubits is $k = \dim C_1 - \dim C_2$, and $y = 0$ recovers an ordinary CSS code. The X-type stabilizers are X^g for $g \in C_2$ and the Z-type stabilizers are $(-1)^{h\cdot y} Z^h$ for $h \in C_1^{\perp}$, so the shift enters only through signs. Codes of this shape are the ones [10] needs for its fault-tolerance machinery, and we call them *CE CSS codes*.

Section IV of [10] settles the small cases and stops at one gap. Their Lemma 4 says that no $[[n, 1, 3]]$ CE CSS code exists for $n \leq 9$. Their Lemma 5 rules out a $[[10, 1, 3]]$ 5-CE one *via their Theorem 3*, i.e. within the dual-rail construction, and they then write, verbatim:

> While this does not rule out the existence of such a code, we are not aware of any construction that achieves it, and its existence remains an open problem.

The sentence is still the current one: as of 2026-08-30 the submission history of arXiv:2507.10395 lists a single version, `[v1]` of 14 July 2025, and neither the abstract page nor the arXiv metadata record carries a journal reference. Their Theorem 6 then exhibits a $[[12, 1, 3]]$ CE CSS code and a $[[14, 3, 3]]$ one, described as the smallest distance-3 CE CSS codes constructed via their Theorem 3, encoding one and multiple logical qubits respectively.

This paper removes the qualification, for $k = 1$ and for $k \geq 2$ alike.

Theorem. For every $n \leq 11$, every $y \in \mathbb{F}_2^n$ and every $k \geq 1$ there is no constant-excitation CSS code on n qubits with k logical qubits and minimum distance at least 3; for $k \geq 2$ there is none for $n \leq 13$ either. Consequently $n = 12$ is the least length at which such a code with one logical qubit exists and $n = 14$ the least length at which one with more than one logical qubit exists; the $[[12, 1, 3]]$ and $[[14, 3, 3]]$ codes of [10] are optimal within the CE CSS class.

Only the range $n \leq 9$ with $k = 1$ was proved before: that case is Lemma 4 of [10], re-derived here by the same exhaustive method as the rest and not extended. The case $n = 10$ is the open problem quoted above. A literature sweep, run on 2026-08-22 and repeated on 2026-08-28 and 2026-08-30, found no *proof* of a minimal-length statement for distance-3 CE CSS codes at either value of k. What it did find is an *assertion* of one: the printed June 2026 snapshot of the Error Correction Zoo [6] says of the $[[12, 1, 3]]$ code, unqualified and citing [10], that "it is the smallest CE CSS code that corrects a single-qubit error" — a statement that entails the nonexistence half of Theorem 5.1, and for which the sweep found no proof. The live zoo entry [5], consulted 2026-08-28 and again 2026-08-30, puts the qualification back, "the smallest distance-three CE CSS code obtainable by dual-rail concatenation", which is what [10] proves. Theorem 5.1 supplies the rest. For $k \geq 2$ we find neither assertion nor refutation in the literature.

The stabilizer question. Constant excitation does not need CSS structure, and without it distance 3 is reached sooner. Two eight-qubit CE stabilizer codes are published: that of Plenio, Vedral and Knight [12], and one of Chang [3]; neither is CSS, and we verify both below. (A $[[10, 1, 3]]$ CE stabilizer code is also on record: [10] notes immediately before its Lemma 4 that dual-rail concatenation of the five-qubit code produces one, and [5] records it under `constant_excitation`.) So the CSS theorem above does not decide the general question — *what is the least n carrying a distance-3 constant-excitation stabilizer code?* — and the two halves of this paper are each other's scope statement: the CSS thresholds are 12 and 14 exactly because the eight-qubit codes are not CSS, and the eight-qubit threshold is compatible with them for the same reason. Nothing here contradicts those codes, and the CSS theorem says about them only that they cannot have the shape (1).

The general question has an assertion in the literature but, so far as we can find, no proof. Plenio, Vedral and Knight write, in [12, §III], verbatim:

> Subsequently a computer search was made for potentially shorter codes; this revealed no such codes, so we conclude that $n = 8$ qubits is the minimum number required for the task of correcting one general error while errors due to the conditional time evolution are corrected perfectly.

What that search covered is not stated — no class, no method, no size — and no description of it has appeared since. The first version of the same preprint (arXiv:quant-ph/9603022v1, 14 March 1996,

titled *Optimal Realistic Quantum Error Correction Code*) claimed instead to *show* the minimality, from a counting bound printed there as Eq. (7),

$$2^l \sum_{i=0}^{t} 3^i \binom{n}{i} \;\le\; \binom{n}{n/2},$$

whose right-hand side is what that version calls "the maximum possible dimension for n encoding qubits" once every codeword is required to carry the same number of excited qubits. Read as $\binom{n}{\lfloor n/2 \rfloor}$ at odd n and taken at $l = t = 1$, it says $2(1 + 3n) \le \binom{n}{\lfloor n/2 \rfloor}$, which first holds at $n = 8$. That argument is unsound twice over: it is a nondegenerate bound, while their own code is degenerate (it contains four weight-2 stabilizers $-Z_i Z_{9-i}$), and it caps the image of the error set by the constant-weight subspace although X- and Y-errors leave that subspace — the corrected nondegenerate count $2(1 + 3n) \le \binom{n}{w-1} + \binom{n}{w} + \binom{n}{w+1}$ already admits $n = 6$. The published version deletes it and puts the sentence quoted above in its place, retaining only the ordinary quantum Hamming bound $2^l \sum_{i=0}^{t} 3^i \binom{n}{i} \le 2^n$ that v1 also had; the deletion is visible in the numbering, and we checked it by compiling both e-print sources: the two logical states are Eqs. (8)–(9) in v1 and Eqs. (7)–(8) in v2, the version accepted by *Phys. Rev. A*. Later work repeats the claim as fact [2] or hedges it [1], which says only that the authors of [12] "had not found a shorter code".

Theorem. For every $n \le 7$, every $y \in \mathbb{F}_2^n$ and every $k \ge 1$ there is no constant-excitation stabilizer code on n qubits with k logical qubits and minimum distance at least 3. Since the code of [12] is such a code with $n = 8$, $k = 1$, $d = 3$, eight is the exact minimum.

We claim the proof, not the statement. The statement for $k = 1$ is in print, asserted without proof in [12] and repeated in [2]; what is new is that it is proved, that it is proved for every $k \ge 1$ at once, and that the non-CSS case — where the codes of [12] and [3] actually live — is reached at all. The CSS sub-case is Lemma 4 of [10], proved there by a constant-weight count that uses the linearity of C_2 and therefore says nothing about non-CSS stabilizer codes. A literature check run on 2026-08-30 found no proof of the general statement: the database of binary stabilizer codes for $n \le 9$ in [4] and the classification of extremal stabilizer codes up to eight qubits in [14] both quotient by local Clifford transformations together with qubit permutations, an equivalence under which constant excitation is not invariant, and neither e-print contains the word "excitation"; the enumeration of inequivalent CSS codes to $n = 14$ in [8] does not contain it either; and we know of no nonexistence result for quantum codes in any proof assistant — the recent stabilizer formalisations such as [9] certify the distance of a given code. What the theorem does *not* cover is codes that are not stabilizer codes, which the sentence of [12] also covers; see Section 12.

Method, and what rests on computation. All three theorems are exhaustive searches, and the paper is mostly about making the searches small enough to run and proving that they lose nothing. Section 2 unfolds (1) into finite combinatorics. Section 3 proves the two structural reductions for the CSS case: constant excitation is *exactly* the containment $C_1 \subseteq V_y$ in a cone V_y determined by y, and the shift may be sorted to $1^w 0^{n-w}$. Section 4 describes the pruned depth-first enumeration and states its soundness; Section 5 gives the $k = 1$ threshold and the size of the search that was run, and verifies the $[[12, 1, 3]]$ code against the same definition the nonexistence theorems refute. Section 6 does the same for $k \ge 2$ and the $[[14, 3, 3]]$ code. Sections 7 and 8 build the stabilizer layer and prove the eight-qubit threshold, with the codes of [12] and [3] verified from their printed equations. Section 9 records the erratum, Section 10 the controls, and Section 11 the machine verification.

Changes in this version. Every theorem of the single-logical-qubit CSS threshold (Sections 2–5) is unchanged in statement and in proof. What changed there is presentation and one count. The twelve-qubit code, a section of its own in the previous version, is now a subsection of Section 5; the paragraph contrasting Definition 2.1 with the bound $\min\{d_1, d_2^{\perp}\}$ of [10, Appendix A.1] has moved from the introduction to Section 2; the $n = 11$ timing was re-measured; and the previous version described the sorted shifts swept at $n = 12$ as twelve of them, in the remark corresponding to Remark 5.3 and in

Proposition 10.1(2). There are thirteen: the search conjoins over $2^w - 1$ for $0 \le w \le 12$. Nothing rests on the count — no theorem states it, and the two sentences that did are corrected here.

New here are: the multi-qubit CSS threshold $n = 14$ (Section 6); the constant-excitation stabilizer layer and the threshold $n = 8$ (Sections 7, 8), together with machine-checked verifications of the eight-qubit codes of [12] and [3]; and the erratum of Section 9. One claim has been scoped rather than extended: where the previous version's abstract called the $[[12, 1, 3]]$ code of [10] optimal *outright* — meaning, as that version's own scope paragraph spelled out, outright rather than only within the dual-rail construction — we now say optimal *within the CE CSS class*, because Theorem 8.1 makes the wider class explicit: distance-3 CE codes exist at eight qubits and are not CSS. No statement of the previous version is withdrawn.

2. Constant-excitation CSS codes as finite combinatorics

Identify $\mathbb{F}_2^n$ with $\{0, 1, \dots, 2^n - 1\}$ by binary expansion, so that coordinate i of v is bit i of v (qubit $i + 1$ of the register), addition in $\mathbb{F}_2^n$ is bitwise XOR $\oplus$, and the coordinatewise product is bitwise AND $\wedge$. Write $\operatorname{wt} v$ for the Hamming weight of v, $\operatorname{supp} v$ for its support, and $\langle u, v\rangle = \operatorname{wt}(u \wedge v) \bmod 2$ for the binary inner product.

Fix n and a shift y. The computational basis kets occurring anywhere in the code space (1) are exactly those indexed by $y + C_1$, since C_1 is the union of the cosets $x^{(j)} + C_2$. Hence the code is $\operatorname{wt}(y)$-CE if and only if $\operatorname{wt}(y \oplus v) = \operatorname{wt}(y)$ for every $v \in C_1$. For $k = 1$ let $\varphi\colon C_1 \to \mathbb{F}_2$ be the nonzero linear functional with $\ker\varphi = C_2$; then $C_1 \setminus C_2 = \{v \in C_1 : \varphi(v) = 1\}$, and since a linear form on C_1 vanishing on the index-two subspace $\ker\varphi$ is 0 or φ,

$$C_2^{\perp} \setminus C_1^{\perp} \;=\; \big\{\, u \in \mathbb{F}_2^n \;:\; \langle u, c\rangle = \varphi(c) \text{ for all } c \in C_1 \,\big\}.$$

The two halves of $d \ge 3$ therefore read as stated in the following definition, which is the object every theorem of Sections 3–5 is about. The passage just made, from (1) to that definition, is an ordinary mathematical argument about states; everything from Definition 2.1 onwards is machine-checked.

Definition 2.1. Let $n \ge 0$ and $y \in \mathbb{F}_2^n$. An $[[n, 1, d \ge 3]]$ *CE CSS code with shift* y is a pair (C_1, φ) consisting of a linear subspace $C_1 \subseteq \mathbb{F}_2^n$ and a map $\varphi\colon C_1 \to \mathbb{F}_2$ additive on C_1, subject to

- (K) $\varphi(v) = 1$ for some $v \in C_1$ (so $k = \dim C_1 - \dim\ker\varphi = 1$);
- (CE) $\operatorname{wt}(y \oplus v) = \operatorname{wt}(y)$ for every $v \in C_1$ (constant excitation, with $w = \operatorname{wt} y$);
- (d_X) $\operatorname{wt}(v) \ge 3$ for every $v \in C_1$ with $\varphi(v) = 1$;
- (d_Z) for every $u \in \mathbb{F}_2^n$ with $1 \le \operatorname{wt}(u) \le 2$ there is $v \in C_1$ with $\langle u, v\rangle \ne \varphi(v)$.

Condition (d_X) says $\min\{\operatorname{wt} v : v \in C_1 \setminus C_2\} \ge 3$ and condition (d_Z) says $\min\{\operatorname{wt} u : u \in C_2^{\perp} \setminus C_1^{\perp}\} \ge 3$; together they say that the stabilizer distance of (1) is at least 3. Appendix A.1 of [10] gives the distance of a CSS code as $\min\{d_1, d_2^{\perp}\}$, where d_1 and $d_2^{\perp}$ are the minimum distances of C_1 and of $C_2^{\perp}$; in general that expression is only a lower bound for the actual distance, and we do not use it. Since Definition 2.1 is the weaker hypothesis, every nonexistence statement below is the stronger one: degenerate codes, whose distance exceeds the bound, are refuted too.

Nothing quantum mechanical survives the translation: Definition 2.1 is a finite object, and for fixed n there are finitely many of them to inspect. Finitely many, but far too many to inspect naively: the number of triples (y, C_1, φ) before any reduction is

$$2^n \sum_r \binom{n}{r}_2 (2^r - 1) \;=\; 8.68 \times 10^{12} \quad (n = 10), \qquad 9.63 \times 10^{14} \quad (n = 11),$$

where $\binom{n}{r}_2$ is the Gaussian binomial coefficient. The next two sections cut this by nine orders of magnitude.

3. The reduction

Everything rests on one identity, which is inclusion–exclusion for supports:

$$\operatorname{wt}(a \oplus b) + 2\,\operatorname{wt}(a \wedge b) \;=\; \operatorname{wt}(a) + \operatorname{wt}(b), \qquad a, b \in \mathbb{F}_2^n. \tag{2}$$

Define the *constant-excitation cone* of a shift y to be

$$V_y \;=\; \big\{\, c \in \mathbb{F}_2^n \;:\; \mathrm{wt}(c) = 2\,|\operatorname{supp} c \cap \operatorname{supp} y| \,\big\}.$$

Theorem 3.1. *Let $y \in \mathbb{F}_2^n$.*

(1) *For a subspace $C_1 \subseteq \mathbb{F}_2^n$, condition* (CE) *holds if and only if $C_1 \subseteq V_y$.*
(2) *If an $[[n, 1, d \geq 3]]$ CE CSS code with shift y exists, then one exists with shift $2^w - 1$, where $w = \mathrm{wt}(y)$; that is, the shift may be taken to be $1^w 0^{n-w}$.*

Proof sketch. (1) Apply (2) with $a = y$, $b = c$: $\mathrm{wt}(y \oplus c) = \mathrm{wt}\, y$ is equivalent to $\mathrm{wt}(c) = 2\,\mathrm{wt}(c \wedge y)$, which is membership in V_y. We stress that this is an equivalence, not merely a necessary condition: the *whole* of (CE) is the containment, so a search that builds C_1 only inside V_y discards nothing.

(2) A permutation of the n coordinates carries a code with shift y to a code with the permuted shift, but it must be shown to act on the entire structure, and this is where the work is. We use bit transpositions: for $i \neq j$ let σ_{ij} swap coordinates i and j, realised as $v \mapsto v$ when v agrees at i and j, and $v \mapsto v \oplus (2^i \,|\, 2^j)$ otherwise. Every clause of Definition 2.1 transports along σ_{ij}: linearity and (K) because σ_{ij} is an $\mathbb{F}_2$-linear involution; (d_X), (d_Z) and (CE) because σ_{ij} preserves Hamming weight and the inner product. Weight invariance is again (2): in the nontrivial case $\sigma_{ij} v = v \oplus m$ with $\mathrm{wt}\, m = 2$ and $\mathrm{wt}(v \wedge m) = 1$.

The iteration is a strong induction on the *numeric value* of y. If some $i < j < n$ has bit i of y clear and bit j set, then $\sigma_{ij} y < y$, so the inductive hypothesis applies to the transported code; if no such pair exists, the set bits of y are downward closed, i.e. $y = 2^{\mathrm{wt}\, y} - 1$ already. Using y itself as the induction measure avoids any bookkeeping over permutations of the coordinate list. □

Part (2) cuts the shifts from 2^n to the $n + 1$ values $2^w - 1$, $0 \leq w \leq n$; part (1) confines C_1 to V_y. A third narrowing concerns φ. Once an ordered list $b_1, \ldots, b_r$ spanning C_1 is fixed, an additive φ is determined by the r values $\varphi(b_1), \ldots, \varphi(b_r)$; conversely every bit mask in $\{0, \ldots, 2^r - 1\}$ arises, and evaluating φ on a linear combination is the $\mathbb{F}_2$ pairing of two masks. This packing is faithful, and is proved to be so; it is what makes "all functionals at this node" a loop over 2^r integers.

4. The search and its completeness

Fix n and a sorted shift $y = 2^w - 1$. The search enumerates the subspaces of V_y depth first. A node holds a basis $b_1, \ldots, b_r$ (most recently added first), the list of pairs (mask, vector) of the subspace S it spans, the last vector added, and, for each error u with $1 \leq \mathrm{wt}\, u \leq 2$, the packed mask of the functional $c \mapsto \langle u, c \rangle$ restricted to S. The last of these is maintained incrementally, one inner product per error per node, and is the reason the search is fast.

At each node the algorithm asks whether *some* nonzero functional on S would make (S, φ) a code. A mask fails if it is zero (violating (K)); or if it is nonzero on some element of S of weight at most 2 (violating (d_X)); or if it coincides with the functional induced by one of the low-weight errors u (violating (d_Z)). The node reports success — in the sense of "no code here" — when every one of the 2^r masks fails.

The node is then extended by a vector $v \in V_y$ subject to three conditions: v exceeds the last vector added; the whole coset $v \oplus S$ lies in V_y; and v is the numerically smallest element of that coset. The second condition is Theorem 3.1(1) applied one step at a time; the third is a canonical-basis condition whose design intent is that each subspace be reached exactly once rather than once per ordered basis. Write Search(n) for the conjunction, over $0 \leq w \leq n$, of the outcome of the whole enumeration at shift $2^w - 1$.

Theorem 4.1. *If* Search$(n) =$ true *then there is no $[[n, 1, d \geq 3]]$ CE CSS code on n qubits, with any shift $y \in \mathbb{F}_2^n$ whatsoever.*

Proof sketch. By Theorem 3.1(2) it suffices to refute codes with sorted shift. Fix such a code (C_1, φ) with shift y. We show that the enumeration visits a node whose span is C_1; since the node reported "no code here", and φ is one of the 2^r masks tested there by faithfulness of the packing, this is a contradiction.

The claim to prove about the pruning is the *at least once* half — that every subspace of V_y is reached — which is the half soundness needs; the "at most once" half is a matter of efficiency and is not used. Suppose the current node spans $S \subsetneq C_1 \subseteq V_y$ and put $v = \min(C_1 \setminus S)$. Then the coset $v \oplus S$ is contained in $C_1 \setminus S$ — if $v \oplus s$ lay in S then so would v, by closure — hence every element of the coset is at least $\min(C_1 \setminus S) = v$, so v is the smallest element of its coset and lies in V_y together with the rest of the coset. It also exceeds the last vector added, by the invariant that every codeword outside the span does. So the branch leading towards C_1 is taken, and the span grows strictly. Formally the recursion carries fuel with the invariant $|C_1| \leq |S| + \text{fuel}$, and the fuel 2^n is never binding. Note that the list of low-weight errors need not be complete for this argument — a shorter list only makes the node test harder to pass — whereas the enumeration of V_y must be complete, and is. □

Remark 4.2. The three narrowings are used, in the formal development, only through Theorems 3.1 and 4.1; no step of the search is assumed sound. This matters here more than usual, because the narrowings shrink the space by nine orders of magnitude, and a search over an accidentally empty space returns "no code" just as loudly as a correct one. Section 10 adds the corresponding computational controls.

5. The threshold for one logical qubit

Theorem 5.1. *For every $n \leq 11$ and every $y \in \mathbb{F}_2^n$ there is no $[[n, 1, d \geq 3]]$ CE CSS code on n qubits with shift y. A $[[12, 1, 3]]$ CE CSS code does exist. Hence twelve is the least length admitting a constant-excitation CSS code with one logical qubit and distance at least* 3*.*

Proof sketch. The positive half is Theorem 5.4 below. The negative half is Theorem 4.1 together with the evaluation of Search(n) for $n = 0, 1, \ldots, 11$; each evaluation returns true. These evaluations are exhaustive finite computations and are, with those of Sections 6 and 8, the only place in the paper where a claim rests on one. Their sizes, counted as (subspace, functional) pairs tested over all sorted shifts, are

n	9	10	11
subspaces of the cones	6 756	35 453	197 484
pairs tested	32 674	225 700	1 504 472

and the smaller n are smaller still; the counts were taken from the auxiliary reimplementation described in the next remark. Every pair is tested; no heuristic cut-off, randomised sampling or partial enumeration enters at any point. The range $n \leq 9$ reproves Lemma 4 of [10]; $n = 10$ answers the question quoted in the introduction; $n = 11$ closes the last case below the known construction. □

Remark 5.2 (timings). Inside the proof assistant the $n = 10$ search takes about 47 seconds and the $n = 11$ search 486–493 seconds of CPU time (three runs; re-measured at 472 seconds on 2026-08-30). The same algorithm written directly in C takes 0.05 and 0.51 seconds respectively; the factor of roughly 950 is the cost of running the evaluator interpreted rather than precompiled, and is not a difference in the algorithm. The C program is what produced the pair counts above, and it agrees with the verified computation on every Boolean outcome. The two are not identical code — the verified one iterates over the cone V_y throughout — so the tabulated counts are the size of the reduced space as the C program measures it. Only the Boolean outcomes are formalised; no theorem depends on the counts.

Remark 5.3 (outside the formal development). The same C program reports 1 390 027 subspaces, 13 637 220 pairs and 46 350 accepting pairs at $n = 12$, spread over the sorted shifts. This computation was *not* carried out inside the proof assistant and nothing above depends on it; we record it only because it bounds how much room the threshold case has. How many inequivalent codes those accepting pairs represent, we do not know.

The twelve-qubit code. The nonexistence theorem is only as meaningful as the definition it refutes, so we check that definition against the code the source actually builds. Appendix A.6 of [10] prints the stabilizer generators

$$g_1 = X_1X_2X_3X_4, \qquad g_2 = X_3X_4X_5X_6, \quad g_3 = X_7X_8X_9X_{10}, \quad g_4 = X_9X_{10}X_{11}X_{12},$$
$$g_5 = Z_1Z_3Z_5Z_7Z_9Z_{11}, \quad g_{5+j} = -Z_{2j-1}Z_{2j} \quad (j = 1, \dots, 6),$$

with logical operators $\bar{X} = X_5X_6X_{11}X_{12}$ and $\bar{Z} = Z_7Z_9Z_{12}$. Reading qubit i as coordinate $i-1$: the four X-type generators span C_2 and $\bar{X}$ extends it to C_1, so

$$C_1 = \operatorname{span}\{15,\ 60,\ 960,\ 3840,\ 3120\}, \qquad \dim C_1 = 5, \quad \dim C_1 - \dim C_2 = 1,$$

and $\varphi = \langle 2368, \cdot\rangle$ is $\bar{Z}$. The shift is pinned by the *signs*: $(-1)^{h\cdot y}Z^h$ must equal $-Z_{2j-1}Z_{2j}$ for each of the six j and $+Z_1Z_3Z_5Z_7Z_9Z_{11}$, which is seven linear conditions on y; all seven hold for $y = 1365$, i.e. $\operatorname{supp} y = \{1, 3, 5, 7, 9, 11\}$ in qubit numbering, and $\operatorname{wt} y = 6$.

Theorem 5.4. *With $n = 12$, $y = 1365$, $C_1 = \operatorname{span}\{15, 60, 960, 3840, 3120\}$ and $\varphi = \langle 2368, \cdot\rangle$, the pair (C_1, φ) is an $[[12, 1, d \geq 3]]$ CE CSS code with shift y in the sense of Definition 2.1, with $|C_1| = 32$. Moreover every ket occurring in its code space has weight exactly 6, so it is 6-CE; every $v \in C_1$ with $\varphi(v) = 1$ has $\operatorname{wt} v \geq 4$ and some such v has $\operatorname{wt} v = 4$; and there is a u of weight 3 with $\langle u, c\rangle = \varphi(c)$ for all $c \in C_1$. Hence $d_X = 4$ and $d_Z = 3$, and the minimum distance is exactly 3.*

Proof sketch. Each clause is a finite check over the 32 elements of C_1, except (d_Z) and the weight-3 witness, which sweep all 4096 vectors of $\mathbb{F}_2^{12}$. All were evaluated exhaustively. □

That $\operatorname{wt} y = 6$ is the value predicted by Theorem 3 of [10] for a dual-rail concatenation of a six-qubit code, which is what this code is; and the value $d_X = 4 > 3 = d_Z$ shows that the two halves of the distance are genuinely different here, so the choice of the true distance rather than the bound in Definition 2.1 is not idle.

6. More than one logical qubit

Definition 2.1 fixes $k = 1$ by choosing a functional. For general k the natural object keeps the two subspaces instead.

Definition 6.1. Let $n, k \geq 0$ and $y \in \mathbb{F}_2^n$. An $[[n, k, d \geq 3]]$ *CE CSS code with shift y* is a pair of linear subspaces $C_2 \subseteq C_1 \subseteq \mathbb{F}_2^n$ with $|C_1| = 2^k|C_2|$ subject to (CE) $\operatorname{wt}(y \oplus v) = \operatorname{wt} y$ for all $v \in C_1$; (d_X) $\operatorname{wt} v \geq 3$ for all $v \in C_1 \setminus C_2$; and (d_Z) for every u with $1 \leq \operatorname{wt} u \leq 2$, if $\langle u, v\rangle = 0$ for all $v \in C_2$ then $\langle u, v\rangle = 0$ for all $v \in C_1$.

For $k = 1$ this is Definition 2.1 with $C_2 = \ker\varphi$, and the passage between the two forms is part of the following proposition.

Proposition 6.2. *Let $C_2 \subseteq C_1$ be an $[[n, k, d \geq 3]]$ CE CSS code with shift y and $k \geq 1$. Enlarging C_2 by one coset inside C_1 produces an $[[n, k-1, d \geq 3]]$ one with the same shift; iterating, an $[[n, 1, d \geq 3]]$ one, which is a code in the sense of Definition 2.1. Consequently Theorem 5.1 already rules out every $[[n, k, d \geq 3]]$ CE CSS code with $k \geq 1$ for $n \leq 11$.*

Proof sketch. Both distance conditions only become easier: $C_1 \setminus C_2$ shrinks, which is (d_X), and $C_2^\perp$ shrinks, which is (d_Z). The index halves at each step by the coset count. So the $k > 1$ question below twelve qubits is free. □

Above eleven qubits it is not free, and the search must be redone, because a code with $k \geq 2$ is a weaker object than a code with $k = 1$ at the same C_1: what it needs is not one surviving functional but a two-dimensional space of them.

Theorem 6.3. *Let $C_2 \subseteq C_1$ be an $[[n, 2, d \geq 3]]$ CE CSS code with shift y. Then there are three maps $\varphi_1, \varphi_2, \varphi_3 \colon C_1 \to \mathbb{F}_2$, each additive on C_1, with $\varphi_3 = \varphi_1 + \varphi_2$ pointwise on C_1, each of them nonzero, each satisfying (d_X) and (d_Z) of Definition 2.1. Moreover the shift of such a datum may be complemented as well as sorted, so it may be taken to be $2^w - 1$ with $2w \leq n$.*

Proof sketch. Choose $x_1 \in C_1 \setminus C_2$ and $x_2 \in C_1 \setminus \langle C_2, x_1\rangle$, which exist because $[C_1 : C_2] = 4$. Then $\langle C_2, x_1\rangle$ and $\langle C_2, x_2\rangle$ are two distinct index-two subgroups between C_2 and C_1; their indicator functionals φ_1, φ_2 and the sum φ_3 are the three nonzero elements of a two-dimensional space of functionals vanishing on C_2, and every such functional inherits (d_X) and (d_Z) from the code. For the complement: only (CE) mentions y, and $\mathrm{wt}(\bar{y} \oplus v) + \mathrm{wt}(y \oplus v) = n$ for all v, so replacing y by its complement $\bar{y}$ preserves (CE) and changes w to $n - w$. Sorting is Theorem 3.1(2) again, transported to this datum. The two together halve the number of shifts to sweep, from $n + 1$ to $\lfloor n/2 \rfloor + 1$. □

The search of Section 4 is reused verbatim; only the node test changes. At a node with span S of dimension r the algorithm now builds the list of *all* surviving masks — those passing the (K), (d_X) and (d_Z) filters — and asks whether three of them close under $\oplus$, since a two-dimensional space of survivors is exactly $\{0, a, b, a \oplus b\}$ with $a, b, a \oplus b$ all surviving. Write $\mathrm{Search}_2(n, y)$ for the outcome of the enumeration at shift y with this node test.

Theorem 6.4. *If* $\mathrm{Search}_2(n, 2^w - 1) = \mathrm{true}$ *for every* w *with* $2w \le n$*, then no* $[[n, k, d \ge 3]]$ *CE CSS code with* $k \ge 2$ *exists on* n *qubits, with any shift.*

Proof sketch. A code with $k \ge 2$ gives one with $k = 2$ by Proposition 6.2 and then the datum of Theorem 6.3, whose shift may be assumed sorted and of weight at most $n/2$. The enumeration half of the argument is verbatim that of Theorem 4.1: the node whose span is C_1 is visited. The terminal half is new. At that node each of the three functionals' packed masks passes all three filters — nonzero because the functional is, orthogonality to the low-weight codeword masks because of (d_X), distinctness from every error-induced mask because of (d_Z) — and the packing respects pointwise $\oplus$, so $\varphi_3 = \varphi_1 + \varphi_2$ becomes exactly the mask triple the node test looks for. No distinctness of the three masks is needed: the three memberships are established independently. □

Theorem 6.5. *For every* $n \le 13$*, every* $k \ge 2$ *and every* $y \in \mathbb{F}_2^n$ *there is no* $[[n, k, d \ge 3]]$ *CE CSS code on* n *qubits with shift* y*. A* $[[14, 3, 3]]$ *CE CSS code does exist, and so does a* $[[14, 2, 3]]$ *one. Hence fourteen is the least length admitting a constant-excitation CSS code with more than one logical qubit and distance at least* 3*, and* $k = 3$ *is already achievable there; the qualification "constructed via Theorem 3" that* [10, §IV] *attaches to its Theorem 6 can be dropped for the multi-qubit code as it was for the single-qubit one.*

Proof sketch. For $n \le 11$ this is Proposition 6.2. For $n = 12$ and $n = 13$ it is Theorem 6.4 together with the evaluation of $\mathrm{Search}_2(n, 2^w - 1)$ at the seven shifts $2w \le n$ in each case; all fourteen evaluations return true. The positive half is Theorem 6.7; a $[[14, 2, 3]]$ code is obtained from it by Proposition 6.2. The auxiliary C reimplementation, run over *all* sorted shifts rather than the half the complement symmetry needs, visits 1 390 027 nodes at $n = 12$ and 9 914 466 at $n = 13$, with 46 350 and 648 900 surviving functionals respectively, of which 6 120 and 85 680 nodes carry three or more, and at most ten survivors at any single node. Inside the proof assistant the shifts for $n \le 12$ together take 227 seconds, the six smaller shifts at $n = 13$ take 203 seconds, and the single shift $y = 63$ at $n = 13$ takes 1 487 seconds of CPU time. □

Remark 6.6 (outside the formal development). The same C program, run in existence mode at $n = 14$ and $k = 2$, exhausts all shifts of weight $0, \ldots, 6$ without finding a code and finds the first one at $y = 127$, of weight 7. So every $[[14, k, 3]]$ CE CSS code with $k \ge 2$ is 7-CE up to complementing the shift. This was not verified inside the proof assistant and nothing above depends on it; the existence half of Theorem 6.5 goes through the explicit witness instead. We note also, as a consistency check between the two searches, that the 46 350 surviving functionals counted at $n = 12$ above is the same number as the accepting pairs of Remark 5.3, as it must be: a surviving functional at a node is an accepting pair for the $k = 1$ search.

The fourteen-qubit code. Appendix A.6 of [10] prints, for its $[[14, 3, 3]]$ code, the X-type generators $g_1 = X_1X_2X_5X_6X_9X_{10}X_{13}X_{14}$, $g_2 = X_3X_4X_5X_6X_{11}X_{12}X_{13}X_{14}$, $g_3 = X_7 \cdots X_{14}$ — masks 13107, 15420, 16320 — together with $g_4 = Z_1Z_3Z_5Z_7Z_9Z_{11}Z_{13}$ (mask 5461, sign +) and seven pair generators

$-Z_{2j-1}Z_{2j}$ (masks $3, 12, 48, 192, 768, 3072, 12288$). The code is obtained there from the dual-rail $[[14, 1, 3]]$ code by deleting two Z-type generators. As before the signs pin the shift: $y \cdot g = 1$ for the seven pairs and $y \cdot g_4 = 0$, satisfied by $y = 5462$, of weight 7.

Theorem 6.7. *Let $n = 14$, $y = 5462$, let C_2 be spanned by $13107, 15420, 16320$ and C_1 by those three together with $15, 51, 195$. Then C_1 is exactly the joint dual of the eight printed Z-type generator masks, $|C_2| = 8$, $|C_1| = 64$, and (C_1, C_2) is a $[[14, 3, d \geq 3]]$ CE CSS code with shift y in the sense of Definition 6.1. Every ket occurring in its code space has weight exactly 7, so it is 7-CE; $d_X = 4$, attained at $v = 15$; and $d_Z = 3$, attained at $u = 21 = Z_1Z_3Z_5$. The minimum distance is exactly 3.*

Proof sketch. Each clause is a finite check: the identification of C_1 with the joint dual sweeps all 2^{14} vectors, (d_Z) sweeps them as well, and the remaining clauses run over the 64 elements of C_1. All were evaluated exhaustively. Pinning C_1 to the dual of the printed Z-generators is what prevents the basis literal from drifting away from the printed code. □

7. Beyond CSS: constant-excitation stabilizer codes

A Pauli operator on n qubits, modulo phases, is a pair $(x, z) \in \mathbb{F}_2^n \times \mathbb{F}_2^n$, acting as $X^x Z^z$. Two Paulis commute if and only if the symplectic form $\langle (x, z), (x', z') \rangle_s = \langle x, z' \rangle + \langle z, x' \rangle$ vanishes on them. A stabilizer group modulo phases is therefore an isotropic subspace $S \subseteq \mathbb{F}_2^{2n}$; the code it defines has $k = n - \dim S$ logical qubits and minimum distance $\min\{\operatorname{wt} P : P \in S^\perp \setminus S\}$, where $\operatorname{wt}(x, z) = |\operatorname{supp} x \cup \operatorname{supp} z|$ and $S^\perp$ is the symplectic dual, i.e. the centraliser. All three of these are properties of the subspace alone; the phases of the generators do not enter. Write $\pi(x, z) = x$ for the projection to the X part, and $Z^{\otimes n} = (0, 1^n)$.

Two facts turn constant excitation into a condition on the subspace. As in Section 2, this passage from the quantum-mechanical definition to a finite object is an ordinary mathematical argument; everything after Definition 7.2 is machine-checked.

Lemma 7.1. *Let S be a stabilizer group on n qubits whose code space is spanned by computational basis kets of one fixed Hamming weight w, and suppose $|b\rangle$ is one of them. Then $Z^{\otimes n} \in S$ modulo phases, and $\pi(S) \subseteq V_b$.*

Proof sketch. $Z^{\otimes n}$ multiplies a weight-w ket by $(-1)^w$, so it acts on the whole code space as the scalar $(-1)^w$; a Pauli acting as a scalar on the code space of a stabilizer code lies in $\pm S$. (This part is Lemma 2(1) of [10], after [7], and is not new.) For the second: the projector $P = |S|^{-1} \sum_{s \in S} s$ sends $|b\rangle$ into the coset $b + \pi(S)$, and the amplitude of $|b \oplus x\rangle$ in $P|b\rangle$, for $x \in \pi(S)$, is a fixed sign times $\sum_{s \in S_Z} \varepsilon(s)(-1)^{z_s \cdot (b \oplus x)}$, where $S_Z = \ker \pi$ is the pure-Z part. Isotropy gives $z_s \cdot x = 0$ for $s \in S_Z$ and $x \in \pi(S)$, so that sum does not depend on x: *a coset that occurs in the code space occurs entirely.* Hence every $b \oplus x$ with $x \in \pi(S)$ has weight $w = \operatorname{wt} b$, which by (2) is exactly $\pi(S) \subseteq V_b$. □

Definition 7.2. Let $n \geq 1$ and $y \in \mathbb{F}_2^n$. A *constant-excitation stabilizer datum* on n qubits with shift y is an $\mathbb{F}_2$-subspace $S \subseteq \mathbb{F}_2^{2n}$ such that

- (I) $\langle P, Q \rangle_s = 0$ for all $P, Q \in S$ (the Paulis commute);
- (K) $2|S| \leq 2^n$, i.e. $\dim S \leq n - 1$, i.e. $k \geq 1$;
- (Z) $Z^{\otimes n} \in S$;
- (CE) $\pi(S) \subseteq V_y$;
- (D) every $P \in S^\perp$ with $1 \leq \operatorname{wt} P \leq 2$ already lies in S.

Condition (D) says that $S^\perp \setminus S$ contains no Pauli of weight ≤ 2, which is the *true* stabilizer distance being at least 3; as in the CSS case, degenerate codes are therefore covered. By Lemma 7.1, an $[[n, k, d]]$ constant-excitation stabilizer code with $k \geq 1$ and $d \geq 3$ yields such a datum, with y any ket occurring in its code space. The converse fails — the cosets of $\pi(S)$ other than the one through y are unconstrained, and no phases are chosen — but that is the direction nonexistence needs: refuting Definition 7.2 for every y refutes every code. Note that (K) is an inequality, so a single statement covers every $k \geq 1$ at once.

Theorem 7.3. *Let $y \in \mathbb{F}_2^n$.*

(1) *If a constant-excitation stabilizer datum with shift y exists, then one exists with shift $2^w - 1$ where $w = \operatorname{wt} y$, and also one with shift $\bar{y}$, of weight $n - \operatorname{wt} y$. Hence it suffices to sweep the shifts $2^w - 1$ with $w \le n/2$.*

(2) *Write $\mathrm{Search}_s(n, y)$ for the outcome of the canonical-basis depth-first enumeration described below at shift y, and $\mathrm{Search}_s(n)$ for its conjunction over the shifts $2^w - 1$, $w \le n/2$. If $\mathrm{Search}_s(n, y) = \mathrm{true}$ then there is no datum with shift y; consequently if $\mathrm{Search}_s(n) = \mathrm{true}$ then there is no constant-excitation stabilizer datum on n qubits with any shift at all, hence no $[[n, k, d]]$ constant-excitation stabilizer code with $k \ge 1$ and $d \ge 3$.*

The enumeration is the one of Section 4, adapted. Pack a Pauli as the integer $v = z + 2^n x$. A new basis vector must exceed the previous one and commute with every basis vector so far — both maintained by filtering the candidate list rather than by re-testing, which is what keeps the cost linear rather than quadratic in the candidate count — must be the numerically least element of the coset it opens and keep that coset's X parts inside V_y, and must satisfy $\pi(v) = 0$ or $Z^{\otimes n} \in \mathrm{span}$. A node reports "this is a code" when its span contains $Z^{\otimes n}$ and every weight-≤ 2 Pauli still commuting with the span is already in it; the search returns the negation.

Proof sketch of Theorem 7.3. (1) A transposition of qubits i and j acts on a packed Pauli by swapping the two X coordinates and the two Z coordinates at once, and all five clauses of Definition 7.2 transport along it, weight and the symplectic form being preserved. Strong induction on the numeric value of y then gives $y = 2^{\operatorname{wt} y} - 1$, exactly as in Theorem 3.1(2). For the complement, $\operatorname{wt}(\bar{v}) + \operatorname{wt}(v) = n$, and only (CE) mentions y.

(2) The last prune is the one that needs an argument, and it is sound *because of the packing*: with $v = z + 2^n x$ every pure-Z Pauli is numerically below every Pauli with a nonzero X part, so by the time the search adds a vector with $\pi(v) \ne 0$ the pure-Z part of S is complete; if $Z^{\otimes n}$ were still outside the span it would be a smaller element of $S \setminus \mathrm{span}$ than the one being added, contradicting minimality. The rest is the argument of Theorem 4.1, with one change: the recursion's fuel is *multiplicative*, the invariant being $|S| \le |\,\mathrm{span}\,| \cdot 2^{\mathrm{fuel}}$ with the span at least doubling at each step. So the fuel counts dimensions, the recursion depth is $n - 1$, and (K) is an invariant of the search rather than a test at the leaves. □

8. Eight qubits

Theorem 8.1. *For every $n \le 7$, every $y \in \mathbb{F}_2^n$ and every $k \ge 1$ there is no constant-excitation stabilizer code on n qubits with k logical qubits and minimum distance at least 3. The code of* [12, Eqs. (7)–(8)] *is such a code with $n = 8$, $k = 1$, $d = 3$. Hence eight is the least length admitting a distance-3 constant-excitation stabilizer code.*

Proof sketch. The negative half is Theorem 7.3 together with the evaluation of $\mathrm{Search}_s(n)$ for $n = 1, \ldots, 7$; every evaluation returns true. This is the only computational input. Its size, as measured by the auxiliary C reimplementation of the same algorithm, is

n	1	2	3	4	5	6	7
shifts	1	2	2	3	3	4	4
nodes	1	8	38	422	3 763	141 784	3 675 228

— 3 821 244 nodes in all, 8.19×10^8 candidate scans and 2.54×10^9 coset checks, with no code found anywhere. No heuristic cut-off, sampling or partial enumeration enters. The positive half is Theorem 8.4. □

Remark 8.2 (timings). The C reimplementation runs the whole range in 32.0 seconds. Inside the proof assistant, with the search kernel precompiled to a shared library, the same evaluations cost 971 seconds of CPU time, of which 862.9 seconds is the single shift $y = 7$ at $n = 7$; the ratio is 27 to 30, which is the cost of the verified evaluator and not a difference in algorithm. Two implementation

choices were measured rather than guessed: filtering the candidate list by "above the last basis vector" as well as by commutation is worth a factor 1.6, and avoiding a power computation in the inner loop is worth 13%.

Remark 8.3. The bound $n \le 7$ is where the structure runs out, and one can see why. For $n \le 7$ the group $\pi(S)$ has order at most 4; order 8 first becomes possible at $n = 8$ with $w = 4$, which is exactly the structure of the code of [12]. This observation comes from an independent search over code spaces described in Section 10 and is not part of the formal development.

The two published eight-qubit codes. Equations (7)–(8) of [12] print the two logical states

$$\begin{aligned}|0_L\rangle &= |00001111\rangle + |11101000\rangle - |10010110\rangle - |01110001\rangle \\ &\quad + |11010100\rangle + |00110011\rangle + |01001101\rangle + |10101010\rangle, \\ |1_L\rangle &= |11110000\rangle - |00010111\rangle + |01101001\rangle - |10001110\rangle \\ &\quad - |00101011\rangle + |11001100\rangle + |10110010\rangle - |01010101\rangle,\end{aligned}$$

(unnormalised), and Eq. (8) of [3] prints a stabilizer directly, namely $ZZXXIIXX$, $XXZZXXII$, $XZXZXIXI$, $-ZIIIZIII$, $-IZIIIZII$, $-IIZIIIZI$, $-IIIZIIIZ$.

Theorem 8.4. (1) *Every one of the sixteen computational basis kets occurring in the two states above has Hamming weight* 4*; the two states have equal norm and are orthogonal, so the code space is two-dimensional. The seven Paulis* $-IIIZZIII$, $-IIZIIZII$, $-IZIIIIZI$, $-ZIIIIIIZ$, $+ZXZIIIXI$, $-YZIYXIIX$, $+XIXZIXIX$ *— each read, here and below, as a sign times* X^xZ^z *with* x *and* z *taken off the letters, so that the one string carrying two* Y*s is the negative of the corresponding product of Pauli matrices — each fix both states with the sign shown, and modulo phases they generate a group* S *of order* 2^7*.*

(2) S *is a constant-excitation stabilizer datum on* 8 *qubits with shift* $1^4 0^4$ *in the sense of Definition 7.2. Its distance is exactly* 3*: the weight-*3 *Pauli* $Z_1Z_3Z_4$ *lies in* $S^\perp \setminus S$*.*

(3) S *is not CSS: its pure-*X *and pure-*Z *elements all lie in an explicit subgroup of order* 16*, so they generate a rank-*4 *subgroup of a rank-*7 *stabilizer.*

(4) *The same holds for the code of* [3, Eq. (8)]*: its seven printed generators generate a group of order* 2^7 *which is a datum with shift* $1^4 0^4$*, of distance exactly* 3 *(witness* $X_1X_5Z_4$*), and not CSS, again with a rank-*4 *pure subgroup.*

Proof sketch. All clauses are finite checks: the group is enumerated (128 elements), the weight, orthogonality and stabilisation checks run over the 256 computational basis kets with integer amplitudes, clause (D) of Definition 7.2 runs over all 4^8 packed Paulis, and non-CSS is refuted by exhibiting the order-16 subgroup and checking that it contains every pure element of S. Here a Pauli (x, z) acts by $X^xZ^z|b\rangle = (-1)^{|z\wedge b|}|b \oplus x\rangle$; every generator above has an even number of Ys, so X^xZ^z differs from the printed string by a real sign and integer arithmetic suffices throughout. □

Part (3) is what makes the two thresholds of this paper compatible: Theorem 5.1 says no $[[8, 1, d \ge 3]]$ CE *CSS* code exists, and both eight-qubit codes are distance-3 CE codes at $n = 8$, so the theorem *predicts* that neither is CSS. That prediction is confirmed independently in part (3), and it could have failed.

Remark 8.5. The authors of [12] anticipated the global shift of (1) by twenty-eight years. Their §III notes that their codewords do not form a linear code — that would force $|00000000\rangle$ to be a codeword and destroy constant excitation — but that the kets of $|0_L\rangle$ form a coset of a linear code with leader $|00001111\rangle$. That leader is exactly the y of (1). What their code has and (1) does not allow are the signs.

9. An erratum

The Error Correction Zoo entry `qubit_8_1_3` [5], consulted on 2026-08-30 both as the live page and as its source file in the zoo's data repository, presents the code of [12] but defines its second logical

state as $|\bar{1}\rangle = X^{\otimes 8}|\bar{0}\rangle$. That is not Eq. (8). The entry's own edit log gives 2026-06-08 as its most recent revision, and the same definition stands verbatim in the corresponding entry of the printed snapshot [6], which is the form of it we cite as fixed.

Theorem 9.1. *Let $|\bar{0}\rangle = |0_L\rangle$ be Eq. (7) of* [12] *and let $|\bar{1}\rangle' = X^{\otimes 8}|\bar{0}\rangle$. Then*

(1) *$|\bar{1}\rangle'$ differs from Eq. (8) in exactly four of the* 256 *amplitudes; the two states have the same eight-ket support, so this is four of its eight nonzero amplitudes, and the four differences are sign flips;*
(2) *the Pauli X_4X_5, of weight* 2*, carries $|\bar{0}\rangle$ to $|\bar{1}\rangle'$ exactly, so the pair $(|\bar{0}\rangle, |\bar{1}\rangle')$ spans a code with a weight-*2 *logical operator: its distance is* 2*, and it is additive;*
(3) *the same operator does* not *carry $|\bar{0}\rangle$ to Eq. (8)'s $|1_L\rangle$.*

Separately, six of the seven Paulis printed in the appendix of the first preprint version of [3] *for the code of* [12] *lie in the stabilizer of Theorem 8.4 with the printed signs, while the seventh, $ZXZIIXII$, anticommutes with two of the other six, so the printed list is not a commuting set; replacing it by $ZXZIIIXI$ repairs this and the list then generates exactly the group of Theorem 8.4.*

Proof sketch. Finite checks over the 256 computational basis kets with integer amplitudes, and membership tests in the 128-element group. □

We state this conservatively and narrowly. It concerns a catalogue entry as consulted on 2026-08-30 and as printed in [6], not the published paper [12], whose Eqs. (7)–(8) are verified in Theorem 8.4 and are unaffected; parts (2) and (3) together say only that the four signs which distinguish Eq. (8) from $X^{\otimes 8}|\bar{0}\rangle$ are what buy the third unit of distance. As of 2026-08-30 we have not contacted the maintainers of [5].

Remark 9.2 (outside the formal development). There is a sharper way to say where the boundary between the two halves of this paper lies, and it is a fact about the published code rather than about a definition. Take the data of [12] itself: $y = |00001111\rangle$, $C_2 = \operatorname{supp}|0_L\rangle \oplus y$ and $C_1 = (\operatorname{supp}|0_L\rangle \cup \operatorname{supp}|1_L\rangle) \oplus y$. Then C_1 is linear of dimension 4, C_2 linear of dimension 3, $C_2 \subset C_1$, and $\operatorname{wt}(y \oplus v) = \operatorname{wt} y = 4$ for every $v \in C_1$: every structural clause of Definition 2.1 holds, with $k = 1$, and $d_Z = 3$. Only (d_X) fails, and it fails at 2. So the code (1) built on PVK's own support, with all coefficients $+1$, is an $[[8, 1, 2]]$ CE CSS code, and their signs are exactly what buys the third unit of distance; Theorem 5.1 says that no *other* choice of (C_1, C_2, y) at $n = 8$ buys it without signs. This computation was carried out in exact integer arithmetic outside the proof assistant, together with its own control: the same pipeline applied to that CSS-by-construction object reports it as CSS, so the CSS test of Theorem 8.4(3) is not a constant "not CSS".

10. Controls

An exhaustive search that returns "nothing found" is worth exactly as much as the evidence that it could have found something. We record the controls; all were checked by the same machinery as the theorems unless the contrary is said.

Proposition 10.1. (1) (The sorting reduction, cross-checked.) *Running the search at $n = 8$ over all* 256 *shifts rather than the nine sorted ones also returns true, giving a second proof of the case $n = 8$ of Theorem 5.1 along a route that never invokes Theorem 3.1(2). The two proofs agree.*

(2) (The reduction, exercised on a real object.) *Transporting the code of Theorem 5.4 along bit transpositions as in the proof of Theorem 3.1(2) yields a code with shift $63 = 2^6 - 1$, which is one of the thirteen sorted shifts $2^w - 1$, $0 \le w \le 12$, that the search visits at $n = 12$; consequently* $\mathrm{Search}(12) = \mathrm{false}$. *The reduced space at the threshold is therefore not accidentally empty.*

(3) (The node test can say yes.) *Take $n = 12$ and shift* 63*. The subspace spanned by* 2145*,* 1105*,* 585*,* 325*,* 195 *lies in V_{63}, and the node test reports there that a functional passing both distance conditions exists.*

(4) (Constant excitation is what closes $n = 10$.) *Let $C_1 \subseteq \mathbb{F}_2^{10}$ be the $[7, 4, 3]$ Hamming code of the Steane code* [13] *padded with three unused coordinates, with φ the parity of the first seven coordinates. Then every clause of Definition 2.1 holds except* (CE) *— it is a genuine $[[10, 1, 3]]$ CSS code — and no shift $y \in \mathbb{F}_2^{10}$ makes it constant excitation. So the case $n = 10$ of Theorem 5.1 is about the constant-excitation hypothesis and not about the CSS or distance conditions.*

(5) (A corrupted code does not slip through.) *Replacing $\bar{Z} = Z_7 Z_9 Z_{12}$ by Z_1 in Theorem 5.4 leaves a nonzero functional all of whose logical-X representatives still have weight at least 3, so (d_X) still holds; but the weight-one vector u supported on qubit 1 then realises the functional, and (d_Z) fails.*

(6) (The definition is the source's.) *Appendix A.1 of* [10] *gives a $[[4, 1, 2]]$ CE CSS code with stabilizers $XXXX$, $-ZZII$, $-IIZZ$ and logical operators $\bar{X} = XXII$, $\bar{Z} = IZZI$. Reading it as above gives $y = 0110$ and $C_1 = \{0000, 0011, 1100, 1111\}$, and the two cosets of* (1) *come out as $\{|0110\rangle, |1001\rangle\}$ and $\{|0101\rangle, |1010\rangle\}$ — the logical basis states printed in* [10] *— with minimum distance exactly 2.*

Item (1) deserves a word. It is a direct computational test of the step that shrinks the shift space from 2^n to $n+1$, and at $n = 8$ it costs 389 354 pairs, about 36 seconds. The same cross-check at $n = 9$ would take about five minutes and at $n = 10$ about eighty, which is why it stops at 8. Item (2) is the complementary test at the other end: it runs the transport on the source's actual code rather than on a hypothetical one. There is also a second, search-free proof of the case $n = 11$ of Theorem 5.1, obtained from the case $n = 10$ by deleting a dead coordinate; the two proofs of that case have the same statement and rest on disjoint computations, one on the 225 700-pair sweep and the other on the 1 504 472-pair one.

Proposition 10.2. *For the multi-qubit search of Section 6:*

(1) (Too strong.) *The statement "no CE CSS code with $k \geq 2$ for $n \leq 14$" is false, by Theorem 6.7; and* $\mathrm{Search}_2(14, 127) = \text{false}$ *follows from that witness through Theorem 6.4 without running any $n = 14$ sweep, which exercises the whole chain of Proposition 6.2 and Theorem 6.3 on a real object.*

(2) (The node test can say yes.) *At the node whose subspace is C_1 of Theorem 6.7, the $k \geq 2$ node test finds its two-dimensional space of surviving functionals.*

(3) (The dimension condition does work.) *At the node of the $[[12, 1, 3]]$ code of Theorem 5.4 — $n = 12$, shift 63, the subspace of Proposition 10.1(3) — exactly ten functionals survive, so the $k = 1$ node test accepts, but no three of them close under $\oplus$, so the $k \geq 2$ test refuses. The $n = 12$ sweep of Theorem 6.5 therefore returns true because of the dimension condition, not because survivors are absent.*

(4) (Corruption.) *Deleting the generator g_3 from C_2 in Theorem 6.7 doubles the index, claiming $k = 4$, and even preserves $d_X \geq 3$; but the weight-one error Z_7 is then orthogonal to C_2 and not to C_1, so $d_Z = 1$. A fourth logical qubit is not one generator away.*

(5) (Non-vacuity, and consistency.) *Weight-2 vectors orthogonal to all of C_2 exist, so clause (d_Z) of Theorem 6.7 quantifies over a nonempty set; and the $k \geq 2$ search returns true on the whole range $n \leq 11$, over all $n + 1$ sorted shifts, where nonexistence is already known through Proposition 6.2.*

Proposition 10.3. *For the stabilizer search of Sections 7–8:*

(1) (The node test can say yes.) *Run on the basis of Theorem 8.4 at $n = 8$, with the low-weight list filtered exactly as the search filters it, the node test returns "this is a code". So the sweeps at $n \leq 7$ return "no code" because there is none, not because the test cannot fire.*

(2) (Too strong.) $\mathrm{Search}_s(8, 1^4 0^4) = \text{false}$, *derived from the datum of Theorem 8.4 through Theorem 7.3(2) with no search run; the verification records that this derivation rests on the eight-qubit data and on no sweep at all.*

(3) (Too weak.) *The Steane* $[[7,1,3]]$ *code* [13]*, packed, is an isotropic subspace of* $\mathbb{F}_2^{14}$ *of order* 64*, so* $k = 1$*, with no weight-*≤ 2 *Pauli in* $S^\perp \setminus S$*, so* $d \geq 3$*: it satisfies every clause of Definition 7.2 except the two constant-excitation ones, and it fails both —* $Z^{\otimes 7} \notin S$*, and no shift* $y \in \mathbb{F}_2^7$ *puts* $\pi(S)$ *inside* V_y*. So Theorem 8.1 is about constant excitation, not about distance-*3 *codes at seven qubits.*

Remark 10.4 (outside the formal development). The nonexistence statement of Theorem 8.1 was also checked by a search written independently, from the code-space side rather than the subspace side: for $k = 1$ exactly two cosets of $\pi(S)$ survive, each a uniform-magnitude superposition with phases in $\{1, i, -1, -i\}$, and constant excitation says both lie in one Hamming weight; the Knill–Laflamme conditions were then tested directly. That search finds no code at $n = 4, 5, 6, 7$, agreeing with Theorem 8.1, and it shares no code and no parametrisation with the enumeration formalised here. Its own controls: a slow full 4^n-Pauli reimplementation agrees at $n = 4, 5, 6$; a differential test of the fast against the slow version at $n = 5, 6, 7$ agrees on all 11 736 configurations tested, including 522 positives; and relaxing to $d \geq 2$ finds 32, 192, 99 120 and 789 376 codes at $n = 4, \ldots, 7$, so the machinery is not vacuous.

11. Verification

Every numbered result of this paper except Lemma 7.1 has been formalised and machine-checked in Lean 4 against Mathlib; there is no `sorry` and no declared axiom anywhere in the development. The exceptions are the two passages from quantum mechanics to finite combinatorics — Lemma 7.1, and the unfolding of (1) into Definition 2.1 in Section 2 — which are ordinary arguments about states in $(\mathbb{C}^2)^{\otimes n}$, and the remarks labelled *outside the formal development*; each is marked where it occurs. Definitions 2.1, 6.1 and 7.2 are structures with the displayed conditions together with the linearity clauses; Theorems 3.1, 4.1, 6.3, 6.4 and 7.3, and all of the lemmas behind them, are ordinary proofs depending only on the standard axioms `propext`, `Classical.choice` and `Quot.sound` — in particular every reduction that shrinks a search space, and every completeness argument for an enumeration, is free of any evaluation axiom. The exhaustive computations — the searches at each $n \leq 11$ for $k = 1$, at $n = 12, 13$ for $k \geq 2$ and at each $n \leq 7$ for stabilizer codes, the unsorted $n = 8$ cross-check, the sweeps verifying the codes of Theorems 5.4, 6.7, 8.4 and 9.1, and the controls of Propositions 10.1–10.3 — are discharged by kernel-external evaluation (`native_decide`), each adding one compiler-trusted axiom for the declaration in which it occurs, and the assembled statements carry exactly the axioms one would predict: four for Theorem 5.1, eighteen for Theorem 6.5, five for the nonexistence half of Theorem 8.1 and fifteen for the whole of it, and nothing else beyond the three standard axioms. Readers who prefer not to trust the compiler can regard those specific finite evaluations as the reproducible computational content of the paper; the C reimplementations mentioned in Sections 5, 6 and 8 agree with all of them, as does the independent code-space search of Section 10. Repository reference to be added at submission.

12. What is left open

Non-stabilizer codes. The sentence of [12] quoted in the introduction covers all codes, and Theorem 8.1 covers stabilizer codes. The remaining case is an optimisation over two-dimensional subspaces of the weight-w subspace of $(\mathbb{C}^2)^{\otimes n}$ — a continuum, not a finite search — so it needs a genuinely different argument, presumably a rank bound on the Knill–Laflamme conditions, rather than more computation.

Classification at the thresholds. The enumerations count nodes and functionals, not codes up to equivalence. At $n = 12$ we do not know how many inequivalent $[[12,1,3]]$ CE CSS codes there are; at $n = 8$ we do not know whether the codes of [12] and [3] are equivalent as stabilizer codes. An exhaustive classification at $n = 8$ in the stabilizer class was priced and not run: extrapolating the growth from $n = 6$ to $n = 7$ puts it at about 10^8 nodes, some fifteen minutes in C but several hours inside the proof assistant.

Larger k at fourteen qubits. Theorem 6.5 settles $k = 2$ and $k = 3$ at $n = 14$ but not $k \geq 4$; the corresponding node test would ask for a three-dimensional space of surviving functionals. Proposition 10.2(4) shows only that the obvious cheat fails.

Distance 4. For $d = 4$ only two thresholds move — $\operatorname{wt} v \geq 4$ in (d_X), and errors of weight up to 3 in (d_Z) and in (D) — and the candidate spaces are unchanged, so the small cases cost the same. We know of no statement of the least n carrying an $[[n, 1, 4]]$ constant-excitation code, CSS or otherwise.

References

[1] G. Alber, Th. Beth, Ch. Charnes, A. Delgado, M. Grassl and M. Mussinger, *Stabilizing distinguishable qubits against spontaneous decay by detected-jump correcting quantum codes*, Phys. Rev. Lett. **86** (2001), no. 19, 4402–4405; arXiv:quant-ph/0103042.

[2] G. Alber, Th. Beth, Ch. Charnes, A. Delgado, M. Grassl and M. Mussinger, *Detected-jump-error-correcting quantum codes, quantum error designs, and quantum computation*, Phys. Rev. A **68** (2003), no. 1, 012316; arXiv:quant-ph/0208140.

[3] E.-J. Chang, *Quantum dual extended Hamming code immune to collective coherent errors*, Phys. Rev. A **112** (2025), no. 5, 052410; arXiv:2503.05249v5. (Preprint versions v1–v4 announced the smallest instance of this family as an $[[8, 1, 4]]$ code; v5 and the published version give $[[8, 1, 3]]$. The appendix of v1 prints a stabilizer for the code of [12]; see Theorem 9.1.)

[4] A. Cross and D. Vandeth, *Small binary stabilizer subsystem codes*, arXiv:2501.17447v1. (A database of one representative of every binary quantum stabilizer code under local Clifford permutation equivalence for $n \leq 9$.)

[5] *The Error Correction Zoo* (V. V. Albert and P. Faist, eds.), entries `constant_excitation`, `css_12_1_3` and `qubit_8_1_3`, `errorcorrectionzoo.org`, consulted 2026-08-28 and 2026-08-30; the last also in its source form, `codes/quantum/qubits/small_distance/small/8/qubit_8_1_3.yml` in the zoo's data repository.

[6] V. V. Albert and P. Faist, *Handbook of Error-Correcting Codes*, arXiv:2606.11484v1 (9 June 2026), entries `css_12_1_3` and `qubit_8_1_3`.

[7] J. Hu, Q. Liang, N. Rengaswamy and R. Calderbank, *Mitigating coherent noise by balancing weight-2 Z-stabilizers*, IEEE Trans. Inform. Theory **68** (2022), no. 3, 1795–1808.

[8] J. M. Koh, A. Gong, A. C. Diaconu, D. B. Tan, A. A. Geim, M. J. Gullans, N. Y. Yao, M. D. Lukin and S. Majidy, *Entangling logical qubits without physical operations*, arXiv:2601.20927v1. (Contains the exhaustive enumeration of all 2.71×10^{10} inequivalent CSS codes up to $n = 14$.)

[9] M. Ehatamm, Y. Lee, X. Wu and R. Tao, *End-to-end formalization of quantum error correction* (the Lean 4 library Lean-QEC, which delivers machine-checked distance certificates for given stabilizer codes), arXiv:2605.16523v1.

[10] C.-Y. Lai, P.-H. Liou and Y. Ouyang, *Fault-tolerant quantum error correction for constant-excitation stabilizer codes under coherent noise*, arXiv:2507.10395v1 (14 July 2025).

[11] Y. Ouyang, *Avoiding coherent errors with rotated concatenated stabilizer codes*, npj Quantum Information **7** (2021), article 87.

[12] M. B. Plenio, V. Vedral and P. L. Knight, *Quantum error correction in the presence of spontaneous emission*, Phys. Rev. A **55** (1997), no. 1, 67–71; arXiv:quant-ph/9603022. (Equation numbers cited here are those of the published paper, which is v2 of the e-print, 21 September 1996. Version v1, 14 March 1996, is titled *Optimal Realistic Quantum Error Correction Code*, carries a counting argument for the minimality as its Eq. (7), and consequently prints the two logical states as Eqs. (8)–(9) rather than (7)–(8); the published version deletes that argument.)

[13] A. M. Steane, *Multiple-particle interference and quantum error correction*, Proc. Roy. Soc. London Ser. A **452** (1996), no. 1954, 2551–2577; arXiv:quant-ph/9601029.

[14] S. Yu, Q. Chen and C. H. Oh, *Graphical quantum error-correcting codes*, arXiv:0709.1780v1. (Classifies all extremal stabilizer codes up to eight qubits.)