CONSTANT-EXCITATION CSS CODES OVER $\mathbb{F}_p$: NINE QUTRITS AND EIGHT QUQUINTS, AGAINST TWELVE QUBITS

KOREA SUPERINTELLIGENCE LABS

Abstract. A quantum code is *constant excitation* (CE) if its whole code space lies in a single eigenspace of the total excitation operator $N = \sum_j \mathrm{diag}(0, \dots, p-1)_j$; such a code acquires only a global phase under the collective coherent drift $e^{-i\theta N}$. For qubits N is the Hamming weight; there Lai, Liou and Ouyang proved that no distance-3 CE CSS code with one logical qubit exists on nine qubits or fewer and exhibited one on twelve, and twelve is now known to be the exact threshold. We ask the same question for p-level systems, which no source we could find asks, and answer it at $p = 3$ and $p = 5$: *there is no* $[[n, 1, d \geq 3]]_3$ *constant-excitation CSS code for any* $n \leq 8$ *and any global shift, and there is one at* $n = 9$*; there is no* $[[n, 1, d \geq 3]]_5$ *one for any* $n \leq 7$*, and there is one at* $n = 8$. So the least block length is 9 for qutrits and 8 for ququints, against 12 for qubits: the qutrit answer is *smaller* than the qubit one, and a table indexed only by $p = 2$ hides that. Both witnesses are exhibited explicitly. We also prove an identity that holds for every CE CSS code over $\mathbb{F}_p$ — the shift's excitation on the support of a codeword v satisfies $2\sum_{i \in \mathrm{supp}\, v} y_i = \mathrm{wt}(v)(p-1)$ — whose $p = 2$ case is the containment condition the qubit treatment takes as a definition. The distance conditions we refute are the true stabilizer distances, so degenerate codes are excluded as well, and only additive closure is assumed of C_1. The nonexistence halves rest on exhaustive searches over a space cut down by three structural theorems; both the reduction and the completeness of the enumeration are theorems, not assumptions, and every statement below is machine-checked in Lean 4.

1. Introduction

Let a register consist of n systems of local dimension p, and let

$$N \;=\; \sum_{j=1}^{n} N_j, \qquad N_j = \mathrm{diag}(0, 1, \dots, p-1)_j,$$

be the total excitation. Write $\mathrm{exc}(v) = \sum_i v_i \in \mathbb{Z}$ for the excitation of a computational basis ket $|v\rangle$, $v \in \mathbb{F}_p^n$ — the digits summed *as integers*, not modulo p. A superposition of basis kets all of the same excitation w is an eigenvector of N with eigenvalue w, so the collective drift $e^{-i\theta N}$ multiplies it by the scalar $e^{-i\theta w}$ for every θ. A code whose entire code space consists of such states is therefore immune to collective coherent noise; following [6] we call it a *w-CE code*. For $p = 2$ one has $N_j = (I - Z_j)/2$, the excitation is the Hamming weight, and this is the familiar qubit notion: the eight-qubit code of Plenio, Vedral and Knight [11] corrects one general error and has all sixteen computational kets of its two logical basis states of weight 4, so it is 4-CE. CE codes have since become the natural carrier for dual-rail concatenation [8], and Hu, Liang, Rengaswamy and Calderbank [4, Corollary 12] proved that a CSS code is oblivious to coherent noise exactly when it is a CE code — which is why the CE class is the right one to index by the local dimension.

Lai, Liou and Ouyang [6, Appendix A.1, Eq. (A1)] obtain CE codes from the CSS construction by adding a *global shift*. Given linear codes $C_2 \subset C_1 \subseteq \mathbb{F}_p^n$ and a vector $y \in \mathbb{F}_p^n$, the logical basis states are

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

Date: August 30, 2026.

where $x^{(j)}$ runs over coset representatives of C_2 in C_1; the number of logical p-ary digits is $k = \dim C_1 - \dim C_2$, and $y = 0$ recovers an ordinary CSS code. We call codes of this shape *CE CSS codes*. The construction is written over $\mathbb{F}_2$ in [6], which is qubit-only from its Section II onwards; reading (1) over $\mathbb{F}_p$ instead is the object of this paper.

Section IV of [6] settles the qubit case up to one gap. Its Lemma 4 shows that no $[[n, 1, 3]]$ CE CSS code exists for $n \le 9$; its Lemma 5 rules out $n = 10$ within their own dual-rail construction and leaves the general case open in as many words; and its Theorem 6 exhibits a $[[12, 1, 3]]$ CE CSS code. The cases $n = 10$ and $n = 11$ were closed in a companion paper [5], so that twelve is the exact qubit threshold within the CE CSS class. That is where the qubit story ends, and it ends at a single value of the local dimension.

The question.

Question 1.1. For each prime p, what is the least n admitting an $[[n, 1, d \ge 3]]_p$ constant-excitation CSS code?

We could not find this question asked anywhere. A literature search carried out on 2026-08-30 recorded the following. All 22 arXiv records matching `all:"constant excitation"` were enumerated individually; six concern quantum codes, five of the six are about qubits, and the sixth is [2]. The only record matching `all:"constant excitation"` together with `all:qudit` is [2], whose constant-excitation spaces are bosonic Fock spaces and whose qudit spaces are permutation-invariant, so it treats neither CE CSS codes over $\mathbb{F}_p$ nor a block-length table indexed by p; `all:"constant excitation"` together with `all:CSS` returns [6] alone. A fresh copy of the Error Correction Zoo data repository [3] has 97 entries under its qudit and Galois-qudit directories and none of the 97 mentions constant excitation, while all 19 CE-mentioning entries are about qubits, bosonic modes or spin-S systems — and the zoo's own CE entry *defines* the property for spin-S systems, so the generalisation is available in print and has not been taken. Searches of the OEIS for the sequence of thresholds returned nothing relevant, which is not merely a coverage gap: the OEIS does carry a sequence of exactly this shape for qubit codes [7]. Minimal block lengths for qudit codes are an active genre — [9, 10] are recent instances — but the codes there are permutation-invariant states of a symmetric subspace rather than constant-excitation CSS codes, and the questions asked are different ones: [9] proves that four qudits are necessary and sufficient for a distance-2 permutation-invariant code encoding one logical qudit of local dimension q, and [10] studies numerically how the block length of a permutation-invariant qudit code scales with its distance. The reader should read all of this as "not found", not as "new".

What is settled here.

Theorem. For every $n \le 8$ and every $y \in \mathbb{F}_3^n$ there is no $[[n, 1, d \ge 3]]_3$ constant-excitation CSS code with shift y, and a $[[9, 1, 3]]_3$ one exists. For every $n \le 7$ and every $y \in \mathbb{F}_5^n$ there is no $[[n, 1, d \ge 3]]_5$ one, and an $[[8, 1, 3]]_5$ one exists. Hence the least block length is 9 at $p = 3$ and 8 at $p = 5$.

These are Theorems 4.1 and 4.2, and both witnesses are printed there. Placed beside the qubit value 12, they give the first three rows of the table the question asks for:

p	2	3	5
least n with an $[[n, 1, d \ge 3]]_p$ CE CSS code	12	9	8

The qutrit entry is *smaller* than the qubit entry. That is the reason the question is not a formality: a table computed at $p = 2$ alone suggests that constant excitation is expensive, and one qutrit more than pays for the constraint that four qubits could not.

We prove one further fact about the objects themselves, elementary but not stated in the source, which never leaves $p = 2$: for every $v \in C_1$ in a CE CSS code with shift y,

$$2 \sum_{i \in \operatorname{supp} v} y_i \;=\; \operatorname{wt}(v)\,(p-1)$$

(Proposition 5.1). At $p = 2$ this reads $2\,|\operatorname{supp} v \cap \operatorname{supp} y| = \operatorname{wt}(v)$, which is exactly the cone membership the qubit treatment installs as a definition, and in particular every codeword of C_1 then has

even weight; over $\mathbb{F}_p$ it is a theorem, and it is the shape any further pruning of the search will have to use.

Scope. Four remarks bound what is being claimed, and each of them makes the nonexistence half *stronger* rather than weaker.

The results are about CSS codes. At $p = 2$, distance-3 CE codes that are not CSS are strictly shorter than twelve qubits: the eight-qubit codes of [11] and of [1] are CE stabilizer codes of distance exactly 3, and neither is CSS — both were recomputed from the authors' own printed data, and in each case the pure-X and pure-Z elements of the stabilizer generate only a rank-4 subgroup of a rank-7 group [5]. The qubit entry 12 is therefore a statement about the CE CSS class, and so is everything in this paper. We say nothing whatever about CE codes over $\mathbb{F}_p$ that are not CSS, and nothing here is evidence about them in either direction.

The distance is the true one. Appendix A.1 of [6] records the distance of a CSS code as $\min\{d_1, d_2^\perp\}$, which in general is only a lower bound for the stabilizer distance. We do not use it. Our conditions are $\min\{\operatorname{wt} v : v \in C_1 \setminus C_2\} \geq 3$ and $\min\{\operatorname{wt} u : u \in C_2^\perp \setminus C_1^\perp\} \geq 3$, so degenerate codes — whose distance exceeds the bound — are refuted too.

Only additive closure is assumed. Definition 2.1 asks C_1 to be closed under addition and φ to be additive, not to be a subspace and a linear form. Over a prime field the two coincide (Lemma 3.1(1)), so nothing is lost, but the nonexistence statements are proved against the formally larger class.

Only $k = 1$ and $d \geq 3$. Everything below fixes one logical p-ary digit and distance at least three. The enumeration of subspaces and its completeness proof do not depend on those choices; the test performed at each node does, and no other cell has been computed.

Plan. Section 2 unfolds (1) over $\mathbb{F}_p$ into finite combinatorics. Section 3 proves the reductions that make the search feasible and states the completeness of the enumeration. Section 4 gives the two thresholds and the sizes of the searches actually run. Section 5 proves the excitation identity. Section 6 records the controls and the $p = 2$ regression. Section 7 describes the machine verification, and Section 8 what is left open, including $p = 7$, where we have no value.

2. CE CSS codes over $\mathbb{F}_p$ as finite combinatorics

Fix a prime p and identify $\mathbb{F}_p^n$ with $\{0, 1, \ldots, p^n - 1\}$ by base-p expansion, so that coordinate i of v is the i-th base-p digit v_i of v, addition in $\mathbb{F}_p^n$ is coordinatewise addition modulo p, and scaling is coordinatewise. Write $\operatorname{wt}(v)$ for the number of nonzero coordinates, $\operatorname{supp} v$ for the set of those coordinates, $\operatorname{exc}(v) = \sum_{i<n} v_i \in \mathbb{Z}$ for the excitation, and $\langle u, v\rangle = \sum_{i<n} u_i v_i \bmod p$ for the $\mathbb{F}_p$ inner product. When a vector is displayed as a digit string, the leftmost digit is coordinate 0.

Fix n and a shift y. The computational basis kets occurring anywhere in (1) are exactly those indexed by $y + C_1$, because C_1 is the union of the cosets $x^{(j)} + C_2$. Hence the code is $\operatorname{exc}(y)$-CE if and only if $\operatorname{exc}(y + v) = \operatorname{exc}(y)$ for every $v \in C_1$. For $k = 1$ let $\varphi\colon C_1 \to \mathbb{F}_p$ be an additive map with $\ker\varphi = C_2$; then $C_1 \setminus C_2 = \{v \in C_1 : \varphi(v) \neq 0\}$, and $C_2^\perp \setminus C_1^\perp$ is the set of $u \in \mathbb{F}_p^n$ with $\langle u, c\rangle = \lambda\varphi(c)$ for all $c \in C_1$ and some $\lambda \neq 0$. This gives the object every theorem below is about.

Definition 2.1. Let p be prime, $n \geq 0$ and $y \in \mathbb{F}_p^n$. An $[[n, 1, d \geq 3]]_p$ *CE CSS code with shift y* is a pair (C_1, φ) consisting of a subset $C_1 \subseteq \mathbb{F}_p^n$ containing 0 and closed under addition, and a map $\varphi\colon C_1 \to \mathbb{F}_p$ with $\varphi(v + w) = \varphi(v) + \varphi(w)$ on C_1, subject to

- (K) $\varphi(v) \neq 0$ for some $v \in C_1$ (so $k = 1$);
- (CE) $\operatorname{exc}(y + v) = \operatorname{exc}(y)$ for every $v \in C_1$ (constant excitation, $w = \operatorname{exc} y$);
- (d_X) $\operatorname{wt}(v) \geq 3$ for every $v \in C_1$ with $\varphi(v) \neq 0$;
- (d_Z) for every $u \in \mathbb{F}_p^n$ with $1 \leq \operatorname{wt}(u) \leq 2$ there is $v \in C_1$ with $\langle u, v\rangle \neq \varphi(v)$.

Condition (d_X) says $\min\{\operatorname{wt} v : v \in C_1 \setminus C_2\} \geq 3$. Condition ($d_Z$) is stated with the single functional φ rather than with all of $\lambda\varphi$, $\lambda \neq 0$; Lemma 3.1(4) shows that the two agree, so (d_Z) really says $\min\{\operatorname{wt} u : u \in C_2^\perp \setminus C_1^\perp\} \geq 3$. Nothing quantum mechanical survives the translation: Definition 2.1 is a finite object.

For each n there are finitely many such objects, and far too many to inspect naively. Before any reduction the number of triples (y, C_1, φ) with C_1 a subspace and $\varphi \neq 0$ is $p^n \sum_r \binom{n}{r}_p (p^r - 1)$, where $\binom{n}{r}_p$ is a Gaussian binomial coefficient; this is 8.8×10^{13} at $(p, n) = (3, 8)$ and 2.2×10^{16} at $(5, 7)$, the largest nonexistence cell of each of the two thresholds below. (This count is elementary arithmetic, recorded here for orientation; no statement in the paper depends on it.) The next section cuts these to 34 040 and 64 207 enumeration nodes respectively.

3. The reduction and the enumeration

Three narrowings make the search feasible, and each is proved rather than assumed. Write

$$V_y \;=\; \{\, c \in \mathbb{F}_p^{\,n} \;:\; \mathrm{exc}(y + c) = \mathrm{exc}(y) \,\}$$

for the *constant-excitation cone* of y, so that (CE) is literally the containment $C_1 \subseteq V_y$.

Lemma 3.1. *Let p be prime and let (C_1, φ) be an $[[n, 1, d \geq 3]]_p$ CE CSS code with shift y.*

1. (Linearity.) *C_1 is closed under scalar multiplication and $\varphi(kv) = k\varphi(v)$ in $\mathbb{F}_p$ for every $v \in C_1$ and every k; that is, additive closure over a prime field is $\mathbb{F}_p$-linearity.*
2. (Transport.) *For $i < j < n$ let σ_{ij} transpose coordinates i and j. Then $(\sigma_{ij} C_1, \varphi \circ \sigma_{ij})$ is an $[[n, 1, d \geq 3]]_p$ CE CSS code with shift $\sigma_{ij} y$.*
3. (Sorting.) *Consequently there is an $[[n, 1, d \geq 3]]_p$ CE CSS code whose shift z has non-increasing digits, $z_0 \geq z_1 \geq \cdots \geq z_{n-1}$. There are $\binom{n+p-1}{p-1}$ such shifts, against p^n in all:* 55 *against* 19 683 *at $(p, n) = (3, 9)$, and* 330 *against* 78 125 *at $(5, 7)$.*
4. (The distance condition is the true one.) *No $u \in \mathbb{F}_p^{\,n}$ with $1 \leq \mathrm{wt}(u) \leq 2$ satisfies $\langle u, \cdot \rangle = \lambda\varphi$ on C_1 for any $\lambda \neq 0$; equivalently,* $\min\{\mathrm{wt}\, u : u \in C_2^{\perp} \setminus C_1^{\perp}\} \geq 3$.

Proof sketch. (1) Iterate additivity: kv is the k-fold sum of v with itself, and likewise for φ. The iteration itself needs nothing of p; what primality buys is the invertibility of scalars, used in (4), in Theorem 3.2 and in Proposition 5.1, and nowhere else in this paper.

(2) One checks that σ_{ij} carries every clause of Definition 2.1 to itself. It is additive, so closure and additivity of φ transport; it permutes coordinates, so it preserves wt, exc and $\langle \cdot, \cdot \rangle$, which handles (CE), (d_X) and (d_Z); and it fixes (K). Each of the three invariances is a reindexing of a sum over $\{0, \ldots, n-1\}$ along the involution σ_{ij}.

(3) Iterate (2) on inversions. The qubit argument — swapping an inversion decreases the numeric value of y — does not survive base p, where a swap need not decrease the value. The measure that does work is $\mu(y) = \sum_{k<n} k\, y_k$: if $i < j$ and $y_i < y_j$ then $\mu(\sigma_{ij} y) < \mu(y)$, because $ib + ja < ia + jb$ whenever $i < j$ and $a < b$. Strong induction on μ terminates at a shift with no inversion, i.e. with non-increasing digits. The count $\binom{n+p-1}{p-1}$ is the number of multisets of size n from p values.

(4) Suppose $\langle u, \cdot \rangle = \lambda\varphi$ on C_1 with $\lambda \neq 0$ and $1 \leq \mathrm{wt}(u) \leq 2$. Since p is prime, λ is invertible; scaling by λ^{-1} preserves weight coordinatewise, so $\lambda^{-1} u$ has the same weight and induces φ itself, contradicting (d_Z). □

The third narrowing concerns φ. Once an ordered list $b_1, \ldots, b_r$ spanning C_1 is fixed, an additive φ is determined by its r values $\varphi(b_1), \ldots, \varphi(b_r)$; by Lemma 3.1(1) additivity already forces φ on every combination. So every additive φ on C_1 is represented by one of the p^r base-p masks in $\{0, \ldots, p^r - 1\}$, and evaluating φ becomes the $\mathbb{F}_p$ pairing of two masks. That this packing is faithful is proved, not assumed; it is what makes "all functionals at this node" a loop over p^r integers.

The enumeration. Fix n and a sorted shift y. The search enumerates the subspaces of V_y depth first. A node carries a basis $b_1, \ldots, b_r$ (most recently added first), the list of (mask, vector) pairs 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 $c \mapsto \langle u, c \rangle$ restricted to S; the last of these is maintained incrementally, one inner product per error per node.

At each node the algorithm asks whether *some* additive φ 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 mask induced by one of the low-weight errors (violating (d_Z)). The node reports "no code here" when all p^r masks fail.

The node is then extended by a vector v subject to three conditions: v exceeds the last vector added; the whole new coset $\bigcup_{\lambda\neq 0}(\lambda v + S)$ lies inside V_y; and v is the numerically least element of that set. The second and third conditions are where $\mathbb{F}_p$ differs from $\mathbb{F}_2$: over $\mathbb{F}_2$ the union is the single coset $v + S$ and one recovers the qubit test. Write $\mathrm{Search}(p, n)$ for the conjunction, over all sorted shifts, of the outcome of the whole enumeration.

Theorem 3.2. *Let p be prime. If* $\mathrm{Search}(p, n) = \mathrm{true}$ *then there is no $[[n, 1, d \geq 3]]_p$ CE CSS code on n systems, with any shift $y \in \mathbb{F}_p^n$ whatsoever.*

Proof sketch. By Lemma 3.1(3) it suffices to refute codes with sorted shift, and the list of shifts the search walks is proved to contain every sorted shift. Fix such a code (C_1, φ) with shift y. We show the enumeration visits a node whose span is C_1; since the node reported "no code here" and φ is one of the p^r masks tested there by faithfulness of the packing, this is a contradiction.

Only the *at least once* half of the pruning is proved, which is the half soundness needs; that each subspace is reached at most once is an efficiency claim and is never used. Suppose the current node spans $S \subsetneq C_1 \subseteq V_y$ and put $v = \min(C_1 \setminus S)$. For $\lambda \neq 0$ and $s \in S$ the element $\lambda v + s$ lies in C_1; it cannot lie in S, for then $\lambda v \in S$ and, λ being invertible *because p is prime*, $v \in S$. So the whole set $\bigcup_{\lambda\neq 0}(\lambda v + S)$ is contained in $C_1 \setminus S \subseteq V_y$ and every element of it is at least $\min(C_1 \setminus S) = v$. Hence v passes the extension test, the branch 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 p^n is never binding. 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. □

For the positive halves we need the converse direction, and it is packaged as a certificate.

Lemma 3.3. *There is a Boolean predicate* $\mathrm{codeCheck}(p, n, y, z, B)$ *of a prime p, a length n, a shift y, a vector z and a list B of vectors, decidable by a finite computation over the span of B and the p^n vectors of $\mathbb{F}_p^n$, such that* $\mathrm{codeCheck}(p, n, y, z, B) = \mathrm{true}$ *implies that* $\big(\mathrm{span}\, B,\ \langle z, \cdot\rangle\big)$ *is an $[[n, 1, d \geq 3]]_p$ CE CSS code with shift y in the sense of Definition 2.1.*

So each witness below is five numbers and a Boolean: p, n, y, the logical-Z vector z, and a basis.

Remark 3.4. The narrowings enter the formal development only through Lemma 3.1 and Theorem 3.2; no step of the search is assumed sound. This matters more than usual here, because the narrowings shrink the space by more than eight orders of magnitude, and a search over an accidentally empty space reports "no code" just as loudly as a correct one. Section 6 supplies the computational controls that this cannot have happened.

4. The thresholds at $p = 3$ and $p = 5$

Theorem 4.1. *For every $n \leq 8$ and every $y \in \mathbb{F}_3^n$ there is no $[[n, 1, d \geq 3]]_3$ constant-excitation CSS code with shift y. A $[[9, 1, 3]]_3$ one does exist: with $n = 9$,*

$$y = 222210000, \qquad C_1 = \mathrm{span}\{\,020200101,\ 022000110,\ 100011000\,\}, \qquad z = 111000000$$

in digit notation (that is, $y = 161$, $C_1 = \langle 7350, 2940, 325\rangle$ and $z = 13$ as base-3 integers, and $\mathrm{exc}(y) = 9$), the pair $(C_1, \langle z, \cdot\rangle)$ satisfies Definition 2.1. Hence ***nine is the least block length admitting a constant-excitation CSS code over $\mathbb{F}_3$ with one logical qutrit and distance at least three****, and it is smaller than the corresponding qubit length* 12*.*

Theorem 4.2. *For every $n \leq 7$ and every $y \in \mathbb{F}_5^n$ there is no $[[n, 1, d \geq 3]]_5$ constant-excitation CSS code with shift y. An $[[8, 1, 3]]_5$ one does exist: with $n = 8$,*

$$y = 44440000, \qquad C_1 = \mathrm{span}\{\,40041001,\ 40401010,\ 44001100\,\}, \qquad z = 22100000$$

in digit notation (that is, $y = 624$, $C_1 = \langle 79254, 16354, 3774\rangle$ and $z = 37$ as base-5 integers, and $\mathrm{exc}(y) = 16$), the pair $(C_1, \langle z, \cdot\rangle)$ satisfies Definition 2.1. Hence ***eight is the least block length***

admitting a constant-excitation CSS code over $\mathbb{F}_5$ with one logical ququint and distance at least three.

Proof sketch of Theorems 4.1 and 4.2. The positive halves are Lemma 3.3 applied to the two displayed certificates; each evaluation of codeCheck is a finite computation, sweeping the 27 (respectively 125) elements of C_1 for (K), (CE) and (d_X), and all $3^9 = 19\,683$ (respectively $5^8 = 390\,625$) vectors of the ambient space for (d_Z).

The negative halves are Theorem 3.2 together with the evaluation of Search$(3, n)$ for $n = 0, \dots, 8$ and of Search$(5, n)$ for $n = 0, \dots, 7$; each evaluation returns true. These evaluations, and the certificate evaluations above, are the only places in this paper where a claim rests on an exhaustive computation. Every node and every mask in the reduced space is visited; no heuristic cut-off, randomised sampling or partial enumeration enters at any point. The sizes of the largest sweeps are

(p, n)	$(3, 7)$	$(3, 8)$	$(5, 5)$	$(5, 6)$	$(5, 7)$
sorted shifts	36	45	126	210	330
errors of weight $1, 2$	98	128	180	264	364
enumeration nodes	4 897	34 040	1 382	8 909	64 207
mask tests	38 991	403 680	9 166	94 005	953 975

and the smaller n are smaller still. As in [5], only the Boolean outcomes are formalised; the node and mask counts were taken from an auxiliary C reimplementation of the same algorithm, which agrees with the verified computation on every Boolean outcome, and no theorem depends on them. □

Remark 4.3 (the exact distances of the two witnesses). Definition 2.1 asks only $d \geq 3$, and that is all the two theorems assert. Outside the formal development, an independent program — digit tuples, brute-force closure of the span, and the two distances computed from $\min \mathrm{wt}(C_1 \setminus C_2)$ and $\min \mathrm{wt}(C_2^{\perp} \setminus C_1^{\perp})$ over the full scalar range — reports for the qutrit witness $|C_1| = 27$, $|C_2| = 9$, full support, $d_X = d_Z = 3$, and for the ququint witness $|C_1| = 125$, $|C_2| = 25$, full support, $d_X = 4$ and $d_Z = 3$. So both have distance exactly 3, and in the ququint case the two halves of the distance differ, which is one more reason not to use the bound $\min\{d_1, d_2^{\perp}\}$ in Definition 2.1. Both witnesses satisfy $\mathrm{exc}(y) = n(p-1)/2$.

Remark 4.4 (timings). With the search compiled to a shared library, the sweeps at $p = 2$ and $p = 3$ together with the qubit and qutrit witnesses and the controls of Section 6 cost 122 seconds of CPU time; the $p = 5$ sweeps and witness cost 354 seconds of CPU time and under 800 megabytes of memory. Measured against a literal C transcription of the same algorithm, the verified evaluation costs about a factor of 4; measured against a table-accelerated C implementation, about a factor of 47. Both figures were taken on a shared twenty-core machine on 2026-08-30 and are properties of the implementations, not of the mathematics.

Remark 4.5 (the cost of constant excitation, outside the formal development). It is natural to compare each threshold with the block length reachable once the constant-excitation clause is dropped. Running the same depth-first search with V_y set to the whole space — *in C only; nothing in this remark is formalised, and no theorem in this paper depends on it* — reports: at $p = 2$, no code at $n = 6$ and a code at $n = 7$, whose C_1 is exactly the $[7, 4, 3]$ Hamming code of the Steane code [12]; at $p = 3$, no code at $n \leq 6$ and a code at $n = 7$; at $p = 5$, no code at $n = 4$ and a code at $n = 5$, whose C_1 is a $[5, 3, 3]_5$ MDS code. At $p = 2$ and $p = 5$ only the cell immediately below the reported value was swept, which is enough: padding C_1 and C_2 with a zero coordinate turns an $[[n, 1, d \geq 3]]_p$ CSS code into an $[[n + 1, 1, d \geq 3]]_p$ one — weights on C_1 are unchanged, and every u in $C_2^{\perp} \setminus C_1^{\perp}$ of the padded code restricts to one of the original of no larger weight — so a single negative cell rules out every smaller length as well. (It is the padding of Proposition 6.1(3), read forwards.) Written beside the thresholds

proved above:

p	2	3	5
least n with a CSS code, CE dropped (C only)	7	7	5
least n with a CE CSS code (proved here)	12	9	8
difference	5	2	3

If the first row is right, the price of constant excitation is not monotone in p. Its $p = 2$ entry is not new — [6] states, in the same Section IV, that "the smallest CSS code that can correct a single-qubit error is the $[[7, 1, 3]]$ Steane code" — and the recovery of that code's C_1 by the unformalised implementation is one more control on it.

5. An identity forced by constant excitation

Constant excitation is stated in Definition 2.1 as one equation per codeword. Because C_1 is closed under scaling (Lemma 3.1(1)), that equation holds simultaneously for the p multiples of a codeword, and averaging over them collapses it to a statement about supports alone. The argument is two lines and we claim no novelty for it; we record it because it is not in the source, which never leaves $p = 2$, and because its $p = 2$ case is the condition the qubit treatment installs as a definition.

Proposition 5.1. *Let p be prime and let (C_1, φ) be an $[[n, 1, d \geq 3]]_p$ CE CSS code with shift y. Then for every $v \in C_1$,*

$$2 \sum_{i \in \operatorname{supp} v} y_i \;=\; \operatorname{wt}(v)\,(p-1).$$

Consequently: for $p = 2$ every $v \in C_1$ has even weight, so C_1 contains no codeword of weight one; and for any p, a codeword of weight one supported at i forces $2y_i = p - 1$.

Proof. By Lemma 3.1(1) each multiple kv, $0 \leq k < p$, lies in C_1, so $\operatorname{exc}(y + kv) = \operatorname{exc}(y)$ for all k and hence $\sum_{k<p} \operatorname{exc}(y + kv) = p \operatorname{exc}(y)$. Expand the left side coordinatewise. A coordinate $i \notin \operatorname{supp} v$ contributes y_i for each k, that is $p\,y_i$. A coordinate $i \in \operatorname{supp} v$ contributes $\sum_{k<p} \big((y_i + kv_i) \bmod p\big)$; since p is prime and $v_i \neq 0$, the map $k \mapsto kv_i$ is a bijection of $\mathbb{F}_p$, so $k \mapsto (y_i + kv_i) \bmod p$ is a bijection too and the contribution is $\sum_{b<p} b = p(p-1)/2$, independently of y_i. Therefore

$$p \operatorname{exc}(y) \;=\; p \sum_{i \notin \operatorname{supp} v} y_i \;+\; \operatorname{wt}(v) \cdot \frac{p(p-1)}{2},$$

and subtracting $p \sum_{i \notin \operatorname{supp} v} y_i$ from $p \operatorname{exc}(y) = p \sum_i y_i$ leaves $p \sum_{i \in \operatorname{supp} v} y_i = \operatorname{wt}(v)\, p(p-1)/2$; now cancel p and multiply by 2. For $p = 2$ the right-hand side of the identity is $\operatorname{wt}(v)$, so $\operatorname{wt}(v)$ is even; a weight-one codeword would make it odd. The last clause is the identity with $\operatorname{wt}(v) = 1$. □

At $p = 2$ the identity reads $2\,|\operatorname{supp} v \cap \operatorname{supp} y| = \operatorname{wt}(v)$, which is the condition defining the cone V_y in the qubit treatment [5]. There it is a definition; here it is a consequence of the excitation identity, and the $p = 2$ case of a statement about every prime.

We stress what is *not* proved. The global form $\operatorname{exc}(y) = |S|(p-1)/2 + \sum_{i \notin S} y_i$, where $S = \operatorname{supp} C_1$ is the union of the supports of the codewords, does not follow from Proposition 5.1: over $\mathbb{F}_2$ the code $\langle 110, 011 \rangle$ has full support but no full-support codeword, so there is no single v to apply the identity to. A proof would need the fibre count $|C_1| = p\,|\{v \in C_1 : v_i = 0\}|$ for $i \in S$, i.e. an average over all of C_1 rather than over the p multiples of one codeword. Section 8 explains why that missing step is the interesting one.

6. Controls, and the regression at $p = 2$

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 checked by the same machinery as the theorems.

Proposition 6.1. (1) (Too strong.) Search$(3, 9)$ = false. *This is derived from the witness of Theorem 4.1 through the contrapositive of Theorem 3.2, not computed afresh; it therefore exercises the soundness theorem against a real object, and it needs the shift $y = 222210000$ to be one of the 55 sorted shifts the search visits, which it is. So the reduced space at the threshold is not accidentally empty.*

(2) (The node test can say yes.) *At $p = 3$, $n = 9$ and shift $y = 222210000$, the node-level test applied to the span of $\{020200101, 022000110, 100011000\}$ reports that a mask passing both distance conditions exists. This is the only computational evidence that the test can ever say yes.*

(3) (Too weak: constant excitation is what closes $n = 8$ at $p = 3$.) *Inside $\mathbb{F}_3^8$ let*

$$C_1 = \mathrm{span}\{\,10000110,\ 20101000,\ 10110000,\ 11000000\,\}, \qquad z = 21110000.$$

Then every clause of Definition 2.1 holds except (CE) *— it is a genuine $[[8, 1, d \geq 3]]_3$ CSS code — and none of the $3^8 = 6561$ global shifts y makes it constant excitation. So the case $n = 8$ of Theorem 4.1 is about the constant-excitation hypothesis and not about the CSS or distance conditions.*

A fourth control needs no separate statement: Definition 2.1 is inhabited at three different primes, by the witnesses of Theorems 4.1 and 4.2 and by the code of Proposition 6.2, so its hypotheses are not vacuously unsatisfiable at any of them.

Item (3) deserves a word: the code there is a $[[7, 1, 3]]_3$ CSS code padded with one unused qutrit, so its support misses one of the eight coordinates. Outside the formal development, the same independent program as in Remark 4.3 reports $|C_1| = 81$, $|C_2| = 27$ and $d_X = d_Z = 3$ for it, so its distance is exactly 3.

The last control is a regression against the qubit case, and it is the strongest evidence that the $\mathbb{F}_p$ machinery computes the intended object: at $p = 2$ it must reproduce results already established by a different implementation, and it does.

Proposition 6.2. *For every $n \leq 11$ and every $y \in \mathbb{F}_2^n$ there is no $[[n, 1, d \geq 3]]_2$ constant-excitation CSS code with shift y; and with $n = 12$, $y = 63$, $C_1 = \mathrm{span}\{2145, 1105, 585, 325, 195\}$ and $z = 2368$ the pair $(C_1, \langle z, \cdot\rangle)$ satisfies Definition 2.1. Hence twelve is the least block length of a CE CSS code at $p = 2$.*

This proposition asserts nothing new. Its range $n \leq 9$ is Lemma 4 of [6]; the cases $n = 10$ and $n = 11$ are the theorem of [5]; and the $n = 12$ code is the code of [6, Appendix A.6], here transported to the sorted shift $y = 63$ and checked against Definition 2.1 rather than against the $\mathbb{F}_2$ definition, so that the two readings of "CE CSS code" are seen to agree on a real object. What is new is only the agreement of two implementations that share no code — the one here is base-p arithmetic by division and remainder, the one in [5] is bitwise — and the agreement is exact at the level of the search and not merely of the verdicts: a C reference implementation of the algorithm formalised here reproduces the published node counts 6756, 35 453 and 197 484 at $n = 9, 10, 11$ and the corresponding counts of tested (subspace, functional) pairs 32 674, 225 700 and 1 504 472, and finds the same $n = 12$ witness after the same 492 711 nodes. (The node test of Section 3 also visits the zero mask, which [5] does not count; its raw mask counts therefore exceed the three just quoted by exactly one per node.) Reproducing a published count with an independent implementation before computing any new number is the discipline these thresholds were computed under.

7. Verification

Every numbered statement in this paper — every definition, lemma, proposition and theorem — has been formalised and machine-checked in Lean 4 against Mathlib. The exceptions are the remarks explicitly labelled as lying outside the formal development (Remarks 4.3 and 4.5, the timings, and the measurements quoted in Section 8); no numbered statement depends on any of them. Definition 2.1

is a structure carrying the four displayed conditions together with additive closure; Lemma 3.1, Theorem 3.2, Lemma 3.3 and all of Proposition 5.1 are ordinary proofs and depend only on the standard axioms `propext`, `Classical.choice` and `Quot.sound` — in particular the entire reduction, the entire completeness argument for the enumeration, and the excitation identity are free of any evaluation axiom. The exhaustive computations — the sweeps Search$(3, n)$ for $n \le 8$, Search$(5, n)$ for $n \le 7$ and Search$(2, n)$ for $n \le 11$, the three certificate evaluations, and the controls of Proposition 6.1(2),(3) — are discharged by kernel-external evaluation (`native_decide`), each contributing one compiler-trusted axiom per sweep to the declarations that use it: Theorem 4.1 carries two such axioms, Theorem 4.2 three, Proposition 6.2 four, and the combined three-prime statement nine. There is no `sorry` and no declared axiom anywhere in the development. Readers who prefer not to trust the compiler can regard those specific finite evaluations as the reproducible computational content of the paper; the independent C implementation mentioned in Section 4 agrees with all of them, and the two witnesses were re-verified by a third program written against the definitions rather than against the search. Repository reference to be added at submission.

8. What is left open

$p = 7$, where we have no value. The obvious next row of the table is not reported here, and deliberately so. Outside the formal development, the C implementation sweeps $n \le 6$ at $p = 7$ in about 16 seconds — 924 sorted shifts, 36 529 nodes, 541 135 mask tests, no code — and the sweep at $n = 7$ did not terminate: it was stopped after nine CPU-minutes here, and an earlier independent run did not finish it in forty CPU-minutes either. A statement that there is no code for $n \le 6$ determines no threshold, so we report no value at $p = 7$, not even a lower bound, which would rest entirely on the unformalised computation. The obstruction is that $|V_y|$ grows with p while the number of sorted shifts grows as $\binom{n+p-1}{p-1}$.

The global excitation identity, which would pay for $p = 7$. If the global form discussed after Proposition 5.1 were proved, together with a deletion lemma saying that a coordinate outside $\operatorname{supp} C_1$ may be removed, the search could be restricted to shifts with $\operatorname{exc}(y) = n(p-1)/2$. Measured in C, that restriction cuts $(p, n) = (3, 8)$ from 34 040 to 21 674 nodes and $(5, 7)$ from 19.2 to 7.6 seconds; both witnesses above already satisfy $\operatorname{exc}(y) = n(p-1)/2$. This is the one missing structural step, and it is what would make $p = 7$ affordable.

Other cells. For $k > 1$ the functional φ becomes a surjection $C_1 \to \mathbb{F}_p^k$ and a node must test the surjective maps rather than the $p^r - 1$ nonzero masks; the enumeration of subspaces and its completeness proof are unchanged. For $d = 4$ only two thresholds move: $\operatorname{wt} v \ge 4$ in (d_X), and errors of weight up to 3 in (d_Z). We know of no statement of the least n carrying an $[[n, 1, 4]]_p$ CE CSS code for any p.

Classification. The enumeration counts accepting pairs, not codes up to equivalence. Outside the formal development, running the $p = 3$, $n = 9$ search to completion over all 55 sorted shifts instead of stopping at the first witness costs 242 401 nodes and finds 5328 accepting (C_1, φ) pairs on 1296 accepting subspaces. How many inequivalent $[[9, 1, 3]]_3$ CE CSS codes those represent we do not know; the relevant group is larger than in the qubit case, since it contains the scalings that preserve the shift as well as the coordinate permutations.

References

[1] 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 that code family as an $[[8, 1, 4]]$ code; v5 and the published version give $[[8, 1, 3]]$.)

[2] A. Aydin, V. V. Albert and A. Barg, *Quantum error correction beyond $SU(2)$: spin, bosonic, and permutation-invariant codes from convex geometry*, PRX Quantum **7** (2026), no. 1, 010341; arXiv:2509.20545.

[3] *Error Correction Zoo*, entry `constant_excitation` and the `qudits` and `qudits_galois` directories of the `eczoo_data` repository, consulted 2026-08-30.

[4] 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; arXiv:2011.00197.

[5] Korea Superintelligence Labs, *No* $[[10, 1, 3]]$ *and no* $[[11, 1, 3]]$ *constant-excitation CSS code: twelve qubits is optimal*, companion paper (2026). Its archival identifier will be supplied at submission, together with that of the present paper.
[6] 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).
[7] *The On-Line Encyclopedia of Integer Sequences*, sequence A393197, "Least number of qubits such that a quantum error correcting code with distance n and one logical qubit exists", consulted 2026-08-30.
[8] Y. Ouyang, *Avoiding coherent errors with rotated concatenated stabilizer codes*, npj Quantum Information **7** (2021), article 87; arXiv:2010.00538.
[9] E. Kubischta and I. Teixeira, *Minimal permutation-invariant qudit codes from edge-colorings of complete graphs*, arXiv:2605.22439v1 (2026).
[10] L. J. Bond, J. Minář, M. Ozols, A. Safavi-Naini and V. Visnevskyi, *Permutation-invariant codes: a numerical study and qudit constructions*, arXiv:2603.10981v1 (2026).
[11] 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.
[12] 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.