THE QUANTUM EXTENSION PROPERTY FOR FULL-SUPPORT STABILIZER CODES: THE MINIMAL LENGTHS AT STABILIZER DIMENSIONS 4 AND 5, DISTANCE-3 COUNTEREXAMPLES, AND A POSITIVE ANSWER FOR CSS CODES

KOREA SUPERINTELLIGENCE LABS

Abstract. The MacWilliams extension theorem fails over the label alphabet $A = \mathbb{F}_2^2$ of qubit Pauli operators, and Mahmoud has recently located the smallest scales at which it fails for *stabilizer* codes. Writing k for the $\mathbb{F}_2$-dimension of the stabilizer, he determines, for $k \leq 3$, the minimal length of a full-support non-extendable weight-preserving isometry (it is 4) and the minimal length of a full-support weight-isometric pair of codes that are not even monomially equivalent (it is 5), and asks for $k \geq 4$. We answer that question at $k = 4$ and $k = 5$: the two minimal lengths are $(4, 5)$ at $k = 4$ and $(5, 6)$ at $k = 5$, so over the four stabilizer dimensions now known the isometry length is $\max(4, k)$, the pair length is $\max(5, k+1)$, and the gap of one between them persists. We answer a second question of his negatively: distance > 2 does not restore the extension property. The perfect $[[5, 1, 3]]$ code carries a weight-preserving automorphism of its stabilizer group implemented by no local Clifford circuit together with a qubit permutation; since every nonzero element of that stabilizer has weight 4, weight preservation is a vacuous hypothesis there, and we therefore also exhibit a pair of $[[6, 1]]$ codes of stabilizer distance at least 3 with a non-extendable weight-preserving isometry whose weight vector takes four distinct values, so that the hypothesis is doing work. A third question we answer positively: for CSS codes with CSS-preserving isometries the extension property holds, at every length, because the class incidence matrix is invertible over $\mathbb{Q}$ — certified by an explicit integral matrix carrying it to a nonzero multiple of the identity, at $k = 3$ and $k = 4$. Every existence statement below is machine-checked against the definitions themselves, the failure of extension being tested against all $6^n\, n!$ monomial transformations, 33 592 320 of them at $n = 6$; the matching lower bounds at $k = 5$ rest on an exhaustive orbit census over 63 082 768 parameterized data and are labelled throughout as computations rather than as formal proofs.

1. Introduction

The objects. Let $A = \mathbb{F}_2^2$ be the label alphabet of a single qubit, carrying the symplectic form

$$\langle (a, b), (a', b') \rangle \;=\; ab' + a'b \qquad \text{in } \mathbb{F}_2,$$

and identify its four elements with Pauli operators modulo phase by $(0, 0) = I$, $(1, 0) = X$, $(0, 1) = Z$, $(1, 1) = Y$; for one qubit this form is the trace-symplectic form and also the trace-Hermitian form on $\mathbb{F}_4$, so no $\mathbb{F}_4$-multiplication occurs anywhere below. On A^n the form is $\langle u, v\rangle = \sum_i \langle u_i, v_i\rangle$, and the weight $\mathrm{wt}(u) = \#\{i : u_i \neq 0\}$ is the number of qubits on which the corresponding Pauli operator acts nontrivially.

An additive (= $\mathbb{F}_2$-linear) subgroup $C \leq A^n$ on which the form vanishes identically is exactly the label set of an abelian stabilizer group, uniquely up to signs [12]; we call it a *stabilizer code*, write $k = \dim_{\mathbb{F}_2} C$ for its *stabilizer dimension*, and record its parameters as $[[n, n-k]]$. The code has *full support* when no coordinate is acted on trivially by all of C. A *monomial transformation* of A^n is a permutation of the n qubits together with an element of $\mathrm{Sp}(2, 2) = \mathrm{GL}(2, 2)$, a group of order 6, in each coordinate; these are precisely the maps induced on labels by local Clifford operations composed with a qubit permutation — each element of $\mathrm{Sp}(2, 2)$ is the label action of a single-qubit Clifford and conversely [7, Rem. 5.3], and two stabilizer codes are monomially equivalent exactly when the corresponding quantum codes are locally-Clifford-permutation equivalent [7, Thm. 5.8] — and there are $6^n\, n!$ of them. Following Mahmoud [1], $\mathrm{QEP}(2, n)$ is the assertion that every weight-preserving additive bijection between stabilizer codes in A^n is the restriction of a monomial transformation —

equivalently, that every weight-preserving isomorphism of n-qubit stabilizer groups is implemented by local Cliffords and a permutation of the qubits.

Classically, MacWilliams' extension theorem [2] says that a Hamming-weight-preserving linear isomorphism between linear codes over a finite field is the restriction of a monomial transformation of the ambient space. The theorem is known to fail over module alphabets with non-cyclic socle [3], which A is; restricting the class of codes weakens the hypothesis, so failure there does not formally imply failure for the self-orthogonal codes — and self-orthogonality is exactly what every stabilizer code has for free. It does not restore the theorem, and Mahmoud [1] determines the smallest scales at which it fails. The quantum question was raised explicitly by Gluesing-Luerssen and Pllaha [6, Question 7.4], who exhibited a four-qubit stabilizer code with a symplectic isometry that is not monomial, and taken up systematically by Pllaha [7]. The contrast with the other MacWilliams theorem is worth naming: the weight-enumerator identities transfer to quantum codes essentially unchanged [13], while the extension theorem does not survive at all.

The source, and the questions it leaves open. Mahmoud [1] proves that $\mathrm{QEP}(2, n)$ holds for $n \leq 2$ and fails at $n = 3$, gives a counterexample family at length $q + 1$ for every qudit dimension q, and — the part this paper continues — analyses the *full-support* question, where the trivial mechanism of an idle qubit is excluded. His Theorem 5.3 states that the minimal length of a full-support counterexample isometry is 4 — its part (i) asserts this over all stabilizer dimensions at once, minimality below $n = 4$ coming from his length-three census, Theorem 4.1 — and that the minimal length of a full-support weight-isometric pair of codes that are not monomially equivalent is 5 *among counterexamples of stabilizer dimension* $k \leq 3$ (his part (ii)); stabilizer dimension $k \leq 2$ admits no full-support counterexample at any length (his Proposition 5.1). Three of the seven questions in his closing section §6.2 are the subject of this paper; we quote them verbatim, and refer to them below by their position in his list. Internal numbers of cited works refer throughout to the arXiv versions recorded in the bibliography.

> **(4) Full-support minimal length as a function of k.** Theorem [5.3] settles $k \leq 3$ (impossible for $k \leq 2$; minimal lengths 4 and 5 for $k = 3$). Determine the minimal full-support lengths for $k \geq 4$, and whether the gap between the minimal isometry-counterexample and the minimal inequivalent pair persists.
>
> **(3) Distance and purity.** Is there a natural "pure/nondegenerate" hypothesis restoring the extension property—e.g., do weight-preserving isometries between stabilizer codes of distance > 2 extend? … Our counterexamples all have distance ≤ 2 or distance 1, so the question is open; note however that stabilizer-element weights ≥ 4 do not suffice to restore extension, by Theorem [5.3](iii).
>
> **(6) Structured subclasses.** Does the extension property hold for natural subclasses of stabilizer codes—CSS codes with CSS-preserving isometries, graph codes, or codes with transitive automorphism group? (For $\mathbb{F}_4$-linear codes and linear isometries it holds; Remark after Example 3.4.)

What is settled here. Question (4) is answered at the next two stabilizer dimensions.

k	minimal isometry length	minimal inequivalent-pair length	source
≤ 2	none at any length	none at any length	[1, Prop. 5.1]
3	4	5	[1, Thm. 5.3]
4	4	5	Theorem 3.1
5	5	6	Theorems 3.2 and 3.3

So the gap of one between the two lengths persists at $k = 4$ and at $k = 5$; over the four dimensions for which the answer is now known the isometry length is $\max(4, k)$ and the pair length is $\max(5, k + 1)$. The structural reason the isometry length stops being 4 and becomes k is that a self-orthogonal code of dimension k needs length at least k: at $k = 3$ that floor is not binding (length 3 is attained but carries no counterexample), and from $k = 4$ on the floor itself is the answer (Proposition 7.2).

Question (3) is answered negatively, twice. The perfect $[[5, 1, 3]]$ code has a weight-preserving automorphism of its stabilizer group that extends to no monomial transformation (Theorem 5.1); this settles the question as posed, but we emphasise that it settles it cheaply, because the perfect code has $A(z) = 1 + 15z^4$, so *every* additive bijection out of it preserves weight and the hypothesis is vacuous. The mechanism is Dyshko's [5, Ex. 3] and Wood's [4], and the closest prior instance is Pllaha's [7, Ex. 4.10], the same phenomenon on a length-5 constant-weight stabilizer code at distance 1; moreover [1, Thm. 5.3(iii)] already records that stabilizer weights ≥ 4 do not restore extension. What Theorem 5.1 adds to that remark is only that such a code can have distance 3. Theorem 5.4 removes the escape: a pair of $[[6, 1]]$ codes of stabilizer distance at least 3, the first of them of distance exactly 3, with a non-extendable weight-preserving isometry whose weight vector takes the four values $2, 4, 5, 6$, so that weight preservation is a genuine constraint and the counterexample is not an instance of the constant-weight mechanism.

Question (6) is answered positively for the CSS subclass (Theorem 6.1): the class incidence matrix of a CSS splitting is invertible over $\mathbb{Q}$, so the weight vector together with the length determines the multiset of coordinate kernels, and every CSS-preserving weight-preserving isometry between CSS stabilizer codes on a common splitting extends to a monomial transformation — at every length. The invertibility is certified at $k = 3$ and $k = 4$ by an explicit integral matrix carrying that incidence matrix to a nonzero multiple of the identity.

How the results are obtained, and what rests on what. All of the positive statements — "there is a counterexample of this shape" — are witnesses tested against the definitions themselves. Non-extendability is the failure of the literal existential over the $6^n\, n!$ monomial transformations; because for a fixed coordinate permutation the n local automorphisms are constrained independently of one another, that existential is decided by running over the $n!$ permutations and searching the six local automorphisms separately in each coordinate — the same predicate, distributed, not a reduction of it. Monomial inequivalence of two *codes* does not factorize that way, and at $n = 6$ it is decided by running over all 33 592 320 maps one at a time. These are machine-checked in Lean 4 (§8) and rest on no lemma of [1]; the only auxiliary observation they use is that a monomial transformation carrying one code onto another carries the generators of the first into the second, which is what licenses refuting monomial equivalence by refuting that necessary condition. The negative statements — "there is no counterexample of this shape below this length" — come in two grades. Those at $k = 4$, $n = 4$ and the floor statements are exhaustive searches that are themselves machine-checked; they use the reduction of a datum to its multiset of coordinate kernels, whose finite content we verify rather than cite (Proposition 7.1), and whose one-line composition into the reduction we state as prose in §2. Those at $k = 5$, $n = 5$ — the lower half of Theorem 3.3 — are an exhaustive orbit census computed outside the formal development, over 1 647 464 data at $n = 5$ and 63 082 768 at $n = 6$; these are recorded as Computations, never as theorems, and §4 states exactly what was run.

2. Preliminaries

Parameterized data. Throughout, $M = \mathbb{F}_2^k$. Following [1, §2.3] we carry a code together with an isometry candidate as a *parameterized datum*: a tuple $\Lambda = (\lambda_1, \ldots, \lambda_n)$ of $\mathbb{F}_2$-linear maps $\lambda_i \colon M \to A$ that are *jointly injective*, $\bigcap_i \ker \lambda_i = 0$, with code

$$C_\Lambda = \{(\lambda_1(x), \ldots, \lambda_n(x)) : x \in M\} \leq A^n,$$

of $\mathbb{F}_2$-dimension exactly k. A pair (Λ, Λ') over the same M encodes the additive bijection $f_{\Lambda,\Lambda'} \colon C_\Lambda \to C_{\Lambda'}$, $\Lambda(x) \mapsto \Lambda'(x)$. Writing $W_\Lambda(x) = \#\{i : \lambda_i(x) \neq 0\}$ for the *weight vector* of Λ, the four conditions in play are, all of them literal:

$$\begin{aligned}
f_{\Lambda,\Lambda'} \text{ preserves weight} &\iff W_\Lambda = W_{\Lambda'};\\
C_\Lambda \text{ is self-orthogonal} &\iff \textstyle\sum_i \langle \lambda_i(x), \lambda_i(y)\rangle = 0 \quad \text{for all } x, y \in M;\\
C_\Lambda \text{ has full support} &\iff \lambda_i \neq 0 \quad \text{for all } i;\\
f_{\Lambda,\Lambda'} \text{ extends monomially} &\iff \exists\, \pi \in S_n,\ \exists\, \tau_i \in \mathrm{GL}(2,2) :\ \lambda'_{\pi(i)} = \tau_{\pi(i)} \circ \lambda_i.
\end{aligned}$$

We call Λ *valid* when C_Λ is self-orthogonal, has full support, and Λ is jointly injective; and we call an ordered pair of valid data at the same k and n with $W_\Lambda = W_{\Lambda'}$ but $f_{\Lambda,\Lambda'}$ not extending monomially a *full-support counterexample isometry.* Two codes $C_\Lambda, C_{\Lambda'}$ are *monomially equivalent* when some monomial transformation carries the first onto the second as a set; a counterexample isometry whose two codes are not monomially equivalent is a *monomially inequivalent pair.* The two minimal lengths of question (4) are the least n at which each of these exists, at a given k.

Coordinate kernels. For $\lambda \neq 0$ the subspace $\ker \lambda \leq M$ has codimension 1 or 2; call $\ker\lambda$ the *class* of the coordinate. Two facts organise every exhaustive search below. First, two nonzero maps $M \to A$ lie in the same $\mathrm{GL}(2,2)$-post-composition orbit if and only if they have the same kernel [1, Lem. 2.5]; combined with the orbit criterion [1, Lem. 2.4] this says that $f_{\Lambda,\Lambda'}$ extends monomially exactly when the two multisets of coordinate kernels coincide, and that C_Λ and $C_{\Lambda'}$ are monomially equivalent exactly when the two multisets lie in one orbit of $\mathrm{GL}(M) = \mathrm{GL}(k,2)$ acting by re-identification of the message space. Second, for λ of rank ≤ 1 the pullback $B_\lambda(x,y) = \langle \lambda x, \lambda y\rangle$ vanishes, the image lying in a line on which an alternating form is zero [1, Lem. 3.1(b)]; and for λ of rank 2, B_λ is the pullback of det along $M \twoheadrightarrow M/\ker\lambda \cong \mathbb{F}_2^2$, so, two rank-2 maps with the same kernel differing by post-composition with $\tau \in \mathrm{GL}(2,2) = \mathrm{SL}(2,2)$ and det being $\mathrm{GL}(2,2)$-invariant, B_λ depends only on $\ker\lambda$. Self-orthogonality of C_Λ is therefore $\sum_i B_{\lambda_i} = 0$ in the $\binom{k}{2}$-dimensional $\mathbb{F}_2$-space of alternating forms on M — an exclusive-or of $\binom{k}{2}$-bit vectors, one per coordinate class. At $k = 3$ this recovers the source's own equation $\sum_{\text{lines } \mathbb{F}_2 v} \eta(\mathbb{F}_2 v)\cdot v = 0$ [1, Lem. 5.2], which is stated there only for $k = 3$.

Consequently a full-support datum is, up to a monomial transformation, a *multiset of n coordinate classes* subject to two conditions: the union of their supports is all of $M \setminus \{0\}$ (joint injectivity) and the exclusive-or of their forms is zero (self-orthogonality). Full support is not a third condition but the requirement that the classes be classes of nonzero maps, which is how the class list is built. There are 14 classes at $k = 3$, 50 at $k = 4$ and 186 at $k = 5$ (Proposition 7.1). The searches of §3 run over these multisets.

Remark 2.1. The passage from data to class multisets is the one place where a step of [1] is used rather than re-proved. Its finite content is verified by exhaustion in Proposition 7.1: that every nonzero map has the kernel of exactly one class representative and is that representative post-composed with an element of $\mathrm{GL}(2,2)$, at each of $k = 3,4,5$; that maps with equal kernels induce equal alternating forms, at $k = 3$ and $k = 4$; and that each $\tau \in \mathrm{GL}(2,2)$ fixes 0, moves nonzero labels to nonzero labels, is additive and preserves the form, so that weight, full support, joint injectivity and self-orthogonality are all invariant under $\lambda_i \mapsto \tau_i \circ \lambda_i$. Only the $k = 4$ instance is consumed below, by Theorem 3.1(iii). What is not formalised is the one-line composition of those facts into the statement "hence every full-support datum is coordinatewise $\mathrm{GL}(2,2)$-equivalent to one built from class representatives, and extension and monomial equivalence of the originals agree with equality and $\mathrm{GL}(k,2)$-equivalence of the class multisets". This affects only the exhaustive negative statement Theorem 3.1(iii); every witness in this paper stands without it.

Distance. For a stabilizer code $C \leq A^n$ the *stabilizer distance* is $d(C) = \min\{\mathrm{wt}(u) : u \in C^\perp \setminus C\}$, the minimum weight of a normalizer element outside the stabilizer — not the minimum weight of C. This is the convention [1] uses: in his §5 the distances of the length-6 pair of his Theorem 5.3(iii) are read off as 1 because their normalizers contain weight-one elements, while every nonzero element of those stabilizers has weight ≥ 4. When C is self-dual, $C^\perp \setminus C$ is empty and the quantity is vacuous; every distance claim below is therefore accompanied by the check $C \subsetneq C^\perp$. The code is *degenerate* when some nonzero element of C has weight $< d(C)$.

3. The minimal full-support lengths at $k = 4$ and $k = 5$

A self-orthogonal $C \leq A^n$ satisfies $C \subseteq C^\perp$ and $|C|\,|C^\perp| = 4^n$, so $|C| \leq 2^n$ and $k \leq n$: the length $n = k$ is a floor. We prove the two instances of that floor we need by exhaustion rather than by citation, so that nothing in this section rests on the counting argument.

Theorem 3.1 ($k = 4$: the minimal lengths are 4 and 5)**.** *Let $k = 4$.*

(i) *There is no self-orthogonal additive code of $\mathbb{F}_2$-dimension 4 in A^n for $n \leq 3$; so no counterexample of stabilizer dimension 4, full support or not, exists below length 4.*

(ii) *At $n = 4$ there is a full-support counterexample isometry. Explicitly, for*

$$C = \langle \mathrm{ZZYZ},\ \mathrm{XZIY},\ \mathrm{IXZZ},\ \mathrm{XXXX} \rangle, \qquad D = \langle \mathrm{YZZY},\ \mathrm{ZIXY},\ \mathrm{IXZZ},\ \mathrm{XXXX} \rangle$$

(each generator list read as the images of a fixed basis of M), both codes are full-support stabilizer codes of dimension 4, the map matching the generators in order preserves weight, and it extends to none of the $6^4 \cdot 4! = 31\,104$ monomial transformations. Hence the minimal full-support counterexample-isometry length at $k = 4$ is exactly 4. These two codes are nevertheless monomially equivalent, so the counterexample is the isometry and not the pair.

(iii) *There is no monomially inequivalent full-support pair of stabilizer dimension 4 at length 4. The search space consists of 4375 valid class multisets; grouping them by weight vector leaves exactly 315 weight classes that contain two distinct multisets, holding 630 multisets in all, and for each of those 315 classes an explicit monomial transformation carries one code onto the other.*

(iv) *At $n = 5$ such a pair does exist:*

$$C = \langle \mathrm{IZZYY},\ \mathrm{XIZIY},\ \mathrm{IXXZZ},\ \mathrm{XXXXX} \rangle, \qquad D = \langle \mathrm{ZIZYZ},\ \mathrm{XIZIY},\ \mathrm{IXXZZ},\ \mathrm{XXXXX} \rangle$$

are full-support $[[5, 1]]$ stabilizer codes of stabilizer dimension 4, the map matching the generators in order preserves weight and extends to none of the $6^5 \cdot 5! = 933\,120$ monomial transformations, and no monomial transformation carries the four generators of C into D, so the two codes are not monomially equivalent.

Consequently the two minimal full-support lengths at $k = 4$ are 4 and 5 — the same pair of values [1] proves for $k = 3$ — and the gap between them persists.

Proof sketch. (i) is an exhaustive search for a chain of four linearly independent, pairwise orthogonal packed labels in A^n for $n = 1, 2, 3$, taken in strictly increasing order and pruned at each node to the labels still orthogonal to every earlier choice; since the form is alternating and bilinear, a subspace is totally isotropic as soon as some basis is pairwise orthogonal, so the search is complete. It returns no chain for $n \leq 3$.

(ii) and (iv) are checks on the displayed data against the definitions of §2: validity, equality of the two weight vectors, and the failure of the existential over π and (τ_i). Nothing is reduced: the existential is the one in §2, decided as described in §1. For the second half of (iv) the test run is the necessary condition that a monomial transformation carrying C_Λ onto $C_{\Lambda'}$ must carry the four generators of C_Λ into $C_{\Lambda'}$; refuting it over all 933 120 maps refutes monomial equivalence.

(iii) is the one exhaustive negative statement. The search runs over multisets of the 50 coordinate classes at $k = 4$ of size 4, keeping those whose supports cover $M \setminus \{0\}$ and whose forms exclusive-or to zero; that this is the whole search space up to monomial equivalence is Remark 2.1. The search returns 4375 multisets, and every multiset it returns is confirmed self-orthogonal by the unreduced definition, so the exclusive-or test is not merely believed. Sorting by weight vector isolates the 315 classes carrying two distinct multisets. The 315 monomial transformations that trivialise them were produced by a separate search and are re-verified here against the image code as a set, so nothing about them is taken on trust. □

Theorem 3.2 ($k = 5$: the minimal isometry length is the floor)**.** *There is no self-orthogonal additive code of $\mathbb{F}_2$-dimension 5 in A^4. At $n = 5$ there is a full-support counterexample isometry of stabilizer dimension 5: for the self-dual $[[5, 0]]$ code*

$$C = \langle \mathrm{YXIIZ},\ \mathrm{IXZYI},\ \mathrm{ZXZIX},\ \mathrm{IXXZZ},\ \mathrm{XXXXX} \rangle$$

the transposition of the first two generators is a weight-preserving automorphism *of C that extends to none of the 933 120 monomial transformations. Hence the minimal full-support counterexample-isometry length at $k = 5$ is exactly 5, the floor $n = k$.*

Proof sketch. The floor is the same exhaustive isotropic-chain search as in Theorem 3.1(i), now for chains of length 5 in A^4, i.e. over the 256 nonzero packed labels; it returns none. The witness is definitional: validity of both presentations, equality of the two weight vectors, equality of the two codes as sets, and the failure of the monomial existential over all 933 120 maps. The weight vector of C takes the values $1, 2, 3, 4, 5$, so weight preservation is not a vacuous hypothesis here. □

Theorem 3.3 ($k = 5$: the minimal inequivalent-pair length is 6)**.** (i) *At $k = 5$, $n = 6$ there is a full-support monomially inequivalent pair. Explicitly,*

$$C = \langle \mathrm{IXZXYI},\ \mathrm{XIIYXZ},\ \mathrm{IXZZXX},\ \mathrm{IXXXZZ},\ \mathrm{IXXXXX} \rangle,$$
$$C' = \langle \mathrm{IIZXZX},\ \mathrm{XXIYIY},\ \mathrm{IXZZXX},\ \mathrm{IXXXZZ},\ \mathrm{IXXXXX} \rangle$$

are full-support $[[6, 1]]$ *stabilizer codes of stabilizer dimension exactly* 5*, the map matching the generators in order preserves weight, that map extends to none of the* $6^6 \cdot 6! = 33\,592\,320$ *monomial transformations, and C and C' are not monomially equivalent.*

(ii) *There is no such pair at length* 5*: below $n = 5$ no stabilizer code of dimension* 5 *exists at all, and at $n = 5$ an exhaustive orbit census over the* 1 647 464 *valid class multisets finds that none of the* 170 128 *counterexample weight classes contains two multisets in different* $\mathrm{GL}(5, 2)$*-orbits.*

Hence the minimal full-support monomially-inequivalent-pair length at stabilizer dimension 5 *is exactly* 6*. With Theorem 3.2, the answer at $k = 5$ is the pair* $(5, 6)$ *against* $(4, 5)$ *at $k = 3$ and $k = 4$: the gap persists.*

Proof sketch. Part (i) is definitional and machine-checked in full. Validity of both data, equality of the weight vectors and the failure of the monomial existential are decided directly. Monomial inequivalence of the two *codes* is the expensive claim, and it is established three ways. (a) No monomial transformation carries the five generators of C into C', checked over all 33 592 320 of them; carrying the code onto the target carries its generators into it, so this refutes equivalence, and it rests on nothing beyond that one sentence. (b) The same search with the per-coordinate labels tabulated, a reorganisation that indexes the local automorphism by the source coordinate rather than the target position; the two agree on a 100-pair sweep at $k = 3$, $n = 4$ and on both pairs of Theorem 3.1. (c) The multiset of codeword supports, up to a relabelling of the coordinates: a monomial transformation is a bijection $C \to C'$ each of whose local factors fixes 0 and moves nonzero labels to nonzero labels, so it carries supports to their permuted images; the two multisets do not agree under any permutation. Outside the formal development the same pair was confirmed inequivalent by brute force over all 33 592 320 monomial transformations at the level of the code sets, by a differently organised pruned search, and by exhausting all 9 999 360 elements of $\mathrm{GL}(5, 2)$ on the class multisets.

Part (ii) is of a different grade and we say so plainly. The statement that no stabilizer code of dimension 5 exists below length 5 is the machine-checked floor of Theorem 3.2: that theorem rules out $n = 4$, and a code in A^n with $n < 4$ becomes one in A^4 on adjoining zero coordinates, which changes neither its dimension nor its self-orthogonality. The statement that nothing at $n = 5$ is monomially inequivalent is an exhaustive computation performed outside the formal development, described in §4: the 1 647 464 valid class multisets at $k = 5$, $n = 5$ are grouped by exact weight vector and the resulting 170 128 counterexample classes are tested for $\mathrm{GL}(5, 2)$-equivalence by union–find under a generating set of $\mathrm{GL}(5, 2)$, with the count 0 reproduced under two different generating sets. It is a measurement, cross-checked, and is not claimed as a formal proof. □

4. The census

The counts quoted in this section were computed outside the formal development, by two independently written programs in two languages, and are reported as measurements. No theorem of this paper depends on them except where Theorem 3.3(ii) says so explicitly.

Computation 4.1. For each stabilizer dimension k and length n below, the number of valid class multisets, the number of distinct weight vectors among them, the number of weight classes holding at least two multisets (the *counterexample classes*), and the number of those that contain two multisets in different $\mathrm{GL}(k, 2)$-orbits (the *monomially inequivalent* ones):

k	n	data	weight vectors	counterexample classes	mon. inequivalent
3	3	63	63	0	0
3	4	350	308	42	0
3	5	1 316	1 030	232	168
3	6	4 102	2 752	883	707
4	4	4 375	4 060	315	0
4	5	58 009	44 549	9 836	630
4	6	524 685	297 460	131 980	75 775
5	5	1 647 464	1 421 784	170 128	0
5	6	63 082 768	43 278 108	14 436 204	937 440

The entry 168 at $k = 3$, $n = 5$ is the number printed in [1, Thm. 5.3(ii)], and it came out of this model before any new cell was computed.

The orbit counts are exact. Canonicalising each multiset over the whole of $\mathrm{GL}(k, 2)$ costs $|\,\mathrm{GL}(k, 2)|$ per item and is hopeless at $k = 5$, where $|\,\mathrm{GL}(5, 2)| = 9\,999\,360$; instead the whole data set is joined by union–find under a generating set of $\mathrm{GL}(k, 2)$, so that two multisets are equivalent exactly when they lie in one component. Two generating sets were used — the $k(k - 1)$ elementary transvections, and the pair consisting of the cyclic coordinate shift together with the single transvection $x \mapsto x + x_0 e_1$ — and the generation was itself verified in the same run by closing the group explicitly and counting 168, 20 160 and 9 999 360 elements at $k = 3, 4, 5$. On every cell where both were run, union–find reproduces the full canonicalisation exactly. The $k = 5$, $n = 6$ row was computed twice, once with each generating set, agreeing in all four columns; grouping by weight vector is a radix sort on a 32-bit hash followed by an exact comparison inside each hash bucket, and 217 061 buckets at that cell held more than one weight vector, so the exact split is load-bearing rather than decorative. The whole census cost about 80 seconds and 1.71 GiB at $k = 5$, $n = 6$, and under 20 seconds per cell elsewhere.

Computation 4.2. Filtering the counterexample classes by stabilizer distance: the number of classes all of whose codes have distance ≥ 3, and how many of those are constant-weight, non-degenerate, or monomially inequivalent.

k	n	distance ≥ 3	constant weight	non-degenerate	mon. inequivalent
3	$5, 6$	0	—	—	—
4	5	1	1	1	0
4	6	0	—	—	—
5	5	13 888	0	(self-dual)	0
5	6	13 516	0	0	0

The single class at $k = 4$, $n = 5$ is the perfect $[[5, 1, 3]]$ code of Theorem 5.1; reproducing it exactly is the control on the distance filter. At $k \leq 3$ nothing of distance ≥ 3 is possible: for an $[[n, n - k, d]]$ code with at least one logical qubit the quantum Singleton bound reads $n - k \leq n - 2d + 2$, that is $k \geq 2(d - 1)$, hence $k \geq 4$ for $d = 3$. That bound is Knill and Laflamme's [10] for qubits, and is proved over a general alphabet, with no purity or nondegeneracy hypothesis, as [9, Thm. 2]; degenerate codes are covered, which is what is needed here, purity entering only in the equality case. The $k = 5$, $n = 5$ row concerns self-dual $[[5, 0]]$ codes, where $C^\perp = C$ and "distance" can only mean the minimum weight of $C \setminus \{0\}$; that is a legitimate invariant of a self-dual additive code but not a code distance, so the row is supporting data and not an answer to anything.

Remark 4.3. The next cell of question (4) is $k = 6$. Counting the valid multisets alone — the cheapest thing that cell can be asked — did not finish in 40 CPU-minutes, and the diagonal grows as 63, 4375, 1 647 464 at $k = 3, 4, 5$, ratios 69 and 377; one more ratio of that order puts $k = 6$, $n = 6$ in the low 10^9, tens of gigabytes for the key array alone. We have not run it.

5. Distance does not restore the extension property

Theorem 5.1 (the perfect five-qubit code)**.** *Let* $C = \langle \mathrm{XZZXI},\ \mathrm{IXZZX},\ \mathrm{XIXZZ},\ \mathrm{ZXIXZ} \rangle$ *be the perfect five-qubit code, a full-support* $[[5, 1]]$ *stabilizer code of stabilizer dimension* 4*. Then* $C \subsetneq C^\perp$ *and*

$d(C) = 3$, and the transposition of the first two generators is a weight-preserving automorphism of C that extends to none of the 933 120 *monomial transformations of A^5. A second, independent witness is the pair of* distinct $[[5,1]]$ *codes*

$$E = \langle \text{ZIZYX},\ \text{XIXZY},\ \text{IZZZZ},\ \text{IXXXX} \rangle, \qquad E' = \langle \text{ZIYXZ},\ \text{XIXZY},\ \text{IZZZZ},\ \text{IXXXX} \rangle,$$

both of distance at least 3*, with a weight-preserving isometry between them extending to no monomial transformation. So the answer to question (3), as posed, is negative: distance > 2 does not restore the extension property.*

Proof sketch. Everything is a check against the definitions. Validity of each datum, equality of the weight vectors, and the failure of the monomial existential are decided over all 933 120 monomial transformations; $C \subsetneq C^{\perp}$ and $d(C) \geq 3$ are decided by running over the $4^5 = 1024$ packed labels of A^5, and $d(C) \geq 4$ is refuted the same way, so the distance of the perfect code is exactly 3. For C and the presentation with its first two generators transposed, the two codes are additionally checked equal as sets, which is what makes the isometry an automorphism. For E and E' what is checked is $d \geq 3$; the test is not vacuous there because $n > k$ forces $C \subsetneq C^{\perp}$, by $|C|\,|C^{\perp}| = 4^n$. □

Remark 5.2 (the hypothesis is vacuous in Theorem 5.1). The perfect five-qubit code has $A(z) = 1+15z^4$: all 15 nonzero stabilizer elements have weight 4. This is machine-checked here, and it is forced: Shor and Laflamme [13] solve the linear-programming constraints coming from their quantum MacWilliams identities for a $[[5,1,3]]$ code and find $A(z) = 1 + 15z^4$, with dual enumerator $B(z) = 1 + 30z^3 + 15z^4 + 18z^5$, to be the only solution — so there is no degenerate five-qubit code, and the witness here is not a hand-picked accident. The same computation is restated in [11, §7.2]. But it means that on this code *every* additive bijection preserves weight, so "preserves weight" is a vacuous hypothesis and the content of Theorem 5.1 is only that the code has a non-monomial additive automorphism. The same is true of E, E', whose weight vectors are also constant. That mechanism is not new: it is used verbatim in [5, Ex. 3] ("all nonzero elements in C_1 have the weight equal to $m-1$ … hence is an isometry"); again in [4], whose length-28 additive code with every nonzero codeword of weight 22 has isometry group $\mathrm{GL}(3, \mathbb{F}_2)$ against a trivial restricted monomial group; and closest of all in [7, Ex. 4.10], a length-5 stabilizer code whose symplectic isometry group is $\mathrm{GL}(3, \mathbb{F}_2)$ against 8 monomial maps. Its seven nonzero elements all have weight 4, and its stabilizer distance is 1; both are read off the generator matrix he displays rather than stated there — the same phenomenon at the same length as Theorem 5.1, but at distance 1. Moreover [1, Thm. 5.3(iii)] already exhibits a counterexample all of whose stabilizer elements have weight ≥ 4 and remarks that such weights do not restore extension. Theorem 5.1 adds to that remark exactly one thing: such a code can have distance 3. Theorem 5.4 is the statement without the escape.

Remark 5.3 (Rains' theorem does not obstruct this). Mahmoud names [8] as the analogue that a positive answer to question (3) would have. It does not apply here, and the reason is a difference of objects, not of hypotheses. Rains takes the automorphism group of an additive code to be, following [12], the subgroup of $S_n \ltimes S_3^n$ preserving the code — monomial by definition — while an *equivalence* is for him an element of $\mathrm{PSU}(2)^{\otimes n}$ composed with a permutation of the qubits. His Corollary 14, that any equivalence of GF(4)-linear quantum codes lies in the Clifford group unless the codes have minimum distance 1 or contain a codeword of weight 2, and his Corollary 16, that every automorphism of a GF(4)-linear code lies in the Clifford group, therefore both constrain *unitary* maps of codespaces. (Corollary 16 is printed with no hypothesis, but its proof appeals to Corollary 14, so it inherits Corollary 14's.) The perfect five-qubit code is GF(4)-linear and meets those hypotheses, so Corollary 16 does govern its codespace; there is no conflict, because the automorphism of Theorem 5.1 is an $\mathbb{F}_2$-linear bijection of the stabilizer group, and such a map need not be induced by any unitary. That gap is what the extension question is about, and [8] does not take it up: weight-preserving isometries and their extension are nowhere discussed there.

Theorem 5.4 (a distance-3 counterexample with a non-vacuous weight hypothesis)**.** *Let*

$$D = \langle \text{ZXYIXY},\ \text{XIXYYX},\ \text{IZZZXY},\ \text{IIIIZZ},\ \text{IXXXXX} \rangle,$$
$$D' = \langle \text{ZYIXXY},\ \text{XIXYYX},\ \text{IZZZXY},\ \text{IIIIZZ},\ \text{IXXXXX} \rangle.$$

Both are full-support $[[6,1]]$ *stabilizer codes of stabilizer dimension exactly* 5 *with* $D \subsetneq D^\perp$, $D' \subsetneq D'^\perp$, $d(D) = 3$ *and* $d(D') \geq 3$. *The map matching the generators in order preserves weight and extends to none of the* $6^6 \cdot 6! = 33\,592\,320$ *monomial transformations of* A^6. *The common weight vector, over the* 32 *points of the message space in order, is*

$$(0, 5, 5, 6, 5, 4, 6, 5, 2, 5, 5, 4, 5, 6, 4, 5, 5, 4, 4, 5, 4, 5, 5, 4, 5, 4, 4, 5, 4, 5, 5, 4),$$

taking the four values $2, 4, 5, 6$ *on the* 31 *nonzero points. So weight preservation is a genuine hypothesis here, and distance* > 2 *does not restore the extension property even when the weight hypothesis is doing work.*

Proof sketch. As before, every clause is decided against the definitions: validity of both data, equality of the weight vectors, the failure of the monomial existential over all $33\,592\,320$ maps, the strict inclusions $D \subsetneq D^\perp$ and $D' \subsetneq D'^\perp$ over the $4^6 = 4096$ packed labels, the two distance bounds, and the displayed weight vector entrywise. The distance test itself was first controlled against the single distance-3 counterexample class of Computation 4.2 at $k = 4$, $n = 5$ — the perfect code of Theorem 5.1 — which it reproduces exactly, and it returns 0 at $k = 3$ and at $k = 4$, $n = 6$. □

Remark 5.5 (what is left open). Two readings of the source's phrase "pure/nondegenerate" come apart here, and the census says so. Every one of the $13\,516$ classes at $k = 5$, $n = 6$ whose codes have distance ≥ 3 is *degenerate*: the value 2 in the weight vector above is a weight-2 stabilizer element, and the scan that additionally demands every nonzero stabilizer element to have weight ≥ 3 returns 0 at that cell (Computation 4.2). So if "pure/nondegenerate" is read as "distance > 2", Theorem 5.4 refutes it, now without the vacuity escape of Remark 5.2; if it is read as "non-degenerate code", this cell does not refute it, and no counterexample computed so far is simultaneously non-degenerate and non-constant-weight. Separately, all $13\,516$ classes have monomially equivalent members, so the *pair*-level version of question (3) — is there a monomially inequivalent weight-isometric pair of distance-3 stabilizer codes? — remains open, now measured empty at $k \leq 5$, $n \leq 6$.

6. CSS codes with CSS-preserving isometries

A CSS stabilizer code splits as $C = C_X \oplus C_Z$ with C_X supported in each coordinate on the line $\{I, X\}$ of A and C_Z on the line $\{I, Z\}$. In the parameterized language this is $M = M_X \oplus M_Z$ with $\dim M_X = k_X$, $\dim M_Z = k_Z$, $k = k_X + k_Z$, and each coordinate map of the form

$$\lambda(u, v) = f(u) \cdot X + g(v) \cdot Z, \qquad (f, g) \in M_X^* \times M_Z^*.$$

A *CSS-preserving* isometry is one respecting the same splitting, i.e. a second datum on the same (k_X, k_Z). The pair (f, g) determines $\ker \lambda = \ker f \oplus \ker g$ and conversely, so there are exactly 2^k coordinate classes, of which $2^k - 1$ are non-idle. Note that CSS coordinates are *not* all of rank one: λ has rank 2 exactly when $f \neq 0$ and $g \neq 0$, which is 9 of the 15 non-idle classes at $k_X = k_Z = 2$ (Proposition 7.2), so self-orthogonality is a genuine constraint on CSS data and the positive answer below is not an artefact of [1, Lem. 3.1(b)].

Let $V = V(k_X, k_Z)$ be the $2^k \times 2^k$ integer matrix whose row indexed by the class c is

$$\big(1,\ [\, x \notin \ker \lambda_c \,]_{x \in M \setminus \{0\}}\big).$$

If η, η' are the multiplicity functions of two data of the same length n on the 2^k classes, then equal lengths say $\sum_c \eta_c = \sum_c \eta'_c$ and equal weight vectors say $\sum_c \eta_c [x \notin \ker \lambda_c] = \sum_c \eta'_c [x \notin \ker \lambda_c]$ for every $x \neq 0$; that is, $\delta^\mathsf{T} V = 0$ for $\delta = \eta - \eta'$.

Theorem 6.1 (question (6) for CSS codes)**.** *For the CSS splittings* $(k_X, k_Z) = (2, 1)$ *and* $(2, 2)$ *— that is, at* $k = 3$ *and* $k = 4$ *— the incidence matrix* $V(k_X, k_Z)$ *carries an explicit integral matrix* Q *with*

$$V \cdot Q = c \cdot I, \qquad c = -4 \quad and \quad c = -256$$

respectively. Hence V *is invertible over* $\mathbb{Q}$, *so a difference of multiplicity functions annihilating every row of* V *vanishes; the multiset of coordinate kernels of a CSS datum is determined by its weight vector together with its length, at* every *length* n, *and every CSS-preserving weight-preserving isometry between*

CSS stabilizer codes on a common splitting extends to a monomial transformation. In the same direction and independently of that argument, the exhaustive search over CSS data finds no counterexample in any computed cell: distinct valid CSS data always have distinct weight vectors, for (k_X, k_Z, n) *equal to* $(2,1,3)$, $(2,1,4)$, $(2,1,5)$, $(2,2,4)$, $(2,2,5)$, $(2,2,6)$, $(1,3,4)$, $(1,3,5)$ *and* $(2,3,5)$*, whose search spaces have sizes* 10*,* 38*,* 101*,* 84*,* 531*,* 2194*,* 147*,* 973 *and* 2016*.*

Proof sketch. The matrix V is built from the model rather than transcribed, and the identity $V \cdot Q = c \cdot I$ is checked entrywise — 64 integer dot products at $k = 3$ and 256 at $k = 4$. What is checked is that identity for the stated c, not that $c = \det V$; the constant does happen to be $\det V$, so that Q is the adjugate, but that was computed outside the development and nothing here uses it. The deduction is two lines and is also *not* part of the formal development: if $\delta^{\mathsf{T}} V = 0$ then $\delta^{\mathsf{T}} V Q = 0$, i.e. $c\,\delta^{\mathsf{T}} = 0$, and $c \neq 0$ gives $\delta = 0$ over $\mathbb{Z}$; the two multiplicity functions therefore agree, the two multisets of coordinate kernels coincide, and the isometry extends by the orbit and kernel criteria [1, Lem. 2.4, 2.5]. The exhaustive cells are the search of §3 restricted to the CSS class list, with the count of weight classes holding two distinct data coming out 0 every time; that the counter is not vacuously zero is Proposition 7.2, which runs it unrestricted at $k = 4$, $n = 4$ and gets $(315, 630)$. □

Remark 6.2 (the equal-length hypothesis is load-bearing). The idle class $(f, g) = (0, 0)$ has identically zero indicator — an idle coordinate contributes to no codeword's weight — so the weight vector alone cannot see how many idle coordinates a datum has. It is exactly the leading all-ones column of V, that is the equal-length condition, that pins the idle multiplicity. On full-support data the idle class is absent and the remaining $2^k - 1$ rows are independent on their own.

Remark 6.3 (the general statement, outside the formal development). The invertibility of $V(k_X, k_Z)$ holds for all k_X, k_Z, by Fourier inversion, and the following argument is offered as mathematics rather than as part of the machine-checked development, which certifies only the two instances of Theorem 6.1. For $x = (u, v) \in M_X \oplus M_Z$,

$$n - W(x) \;=\; \sum_{(f,g)} \eta(f,g)\,[\,f(u) = 0\,]\,[\,g(v) = 0\,].$$

The 2^{k_X} functions $u \mapsto [f(u) = 0] = \frac{1}{2}(1 + \chi_f(u))$ are the image of the character basis $\{\chi_f\}$ under an invertible transform, hence a basis of the functions on M_X; likewise on M_Z; hence their 2^k products form a basis of the functions on $M_X \times M_Z$. So η is determined by $x \mapsto n - W(x)$, that is by the weight vector together with the length, and therefore so is the multiset of coordinate kernels. Equal weight vectors and equal lengths thus force the isometry to extend, for every k_X, k_Z and every n; no hypothesis of self-orthogonality or of full support is used. Formalising this is the obvious next step and would subsume Theorem 6.1.

7. Controls

A development in which every predicate is vacuously true, or in which no search ever succeeds, proves nothing. The following two propositions are what makes the statements above readable as claims; both are machine-checked.

Proposition 7.1 (the source's own census, reproduced)**.** *Rebuilt from the definitions of §2 and run before any new cell was computed, the search of §3 returns at* $k = 3$ *exactly* 63*,* 350 *and* 1316 *valid class multisets at lengths* 3*,* 4 *and* 5*, with* 0*,* 42 *and* 232 *counterexample weight classes holding* 0*,* 84 *and* 518 *multisets — the numbers behind* [1, Thm. 5.3(i),(ii)]*, so that the first counterexample is at* $n = 4$ *and the first inequivalent pair at* $n = 5$*. Outside the formal development the same program returns exactly* 168 $\mathrm{GL}(3,2)$*-classes containing monomially inequivalent pairs at* $n = 5$*, which is the count printed in* [1, Thm. 5.3(ii)]*. Moreover there are exactly* 14*,* 50 *and* 186 *coordinate-kernel classes at* $k = 3, 4, 5$*; every nonzero map* $M \to A$ *has the kernel of exactly one class representative and is that representative post-composed with an element of* $\mathrm{GL}(2,2)$*, at each of* $k = 3, 4, 5$*; maps with equal kernels induce equal alternating forms, at* $k = 3$ *and* $k = 4$*; and every multiset the reduced search returns at* $k = 4$*,* $n = 4$ *is self-orthogonal by the unreduced definition.*

Proposition 7.2 (negative controls)**.** *The following are all machine-checked. The model: the form is alternating and bilinear on A; each of the six elements of* $\mathrm{GL}(2,2)$ *fixes* 0*, moves nonzero labels to nonzero labels, is additive and preserves the form; and those six are* all *the additive bijections of A fixing* 0*, checked over all* 256 *tables, so the monomial group is not undercounted. Too-large claims are refuted: the extension test and the code-equivalence test both return true on a genuine monomial image, so "nothing extends" is false and weight preservation alone does not defeat extension. Too-small claims are refuted: a self-orthogonal code of dimension k at length k exists for $k = 3, 4, 5$, so the floor statements are not vacuous emptiness, and at $k = 3$ the floor is* not *binding — length* 3 *is attained and carries no counterexample — which is why the answer there is* 4 *and not* 3*. The three conditions in validity are independent: four copies of one rank-one map are self-orthogonal and full-support but not jointly injective; a datum with a zero coordinate loses full support; and a datum with one rank-two coordinate fails self-orthogonality where the all-rank-one datum satisfies it, which is also a check on* [1, Lem. 3.1(b)]*. Distance: the distance test is vacuously true on a self-dual code, so the strict inclusion $C \subsetneq C^{\perp}$ is checked wherever a distance is claimed, and a code is exhibited on which the distance test genuinely fails. CSS:* 9 *of the* 15 *non-idle CSS classes at $k_X = k_Z = 2$ have rank two (*3 *of* 7 *at* $(2,1)$*,* 7 *of* 15 *at* $(1,3)$*); the CSS data really are valid stabilizer data; and the identical counting code run over the unrestricted class list at $k = 4$, $n = 4$ returns* $(315, 630)$*, so the zero of Theorem 6.1 is a fact about CSS codes and not about the counter.*

8. Verification

Every statement above that is called a theorem or a proposition has been formally verified in Lean 4 [14] (toolchain `v4.32.0`), with exactly two exceptions, both marked in place as computed outside it: the second half of Theorem 3.3(ii), which is the census of §4, and the sentence of Proposition 7.1 recording the orbit count 168. The repository is private, repository reference to be added at submission. The development contains no `sorry`, and its modules import neither *Mathlib* [15] nor any other library — every object is a natural number or a list, and the definitions of §2 are transcribed directly, so that non-extendability really is decided by quantifying over coordinate permutations and $\mathrm{GL}(2,2)$-elements rather than through any criterion. The finite checks are carried out by Lean's compiled evaluator (`native_decide`), so each theorem depends on `propext` together with one evaluation axiom per conjunct, and on nothing else: no `sorryAx`, no `Classical.choice`, no `Quot.sound`. The model facts of Proposition 7.2 concerning the form and $\mathrm{GL}(2,2)$ are the exception, being small enough to be decided by the kernel itself with `propext` alone. Trusting the compiled evaluator is the price of the 33 592 320-map searches; it is the only trusted step, and the searches most exposed to it were re-run outside Lean by two further programs sharing no code, as §4 and the proof of Theorem 3.3 record. Three things stated above sit outside the verification and are flagged where they occur: the composition step of the kernel reduction (Remark 2.1), which affects only Theorem 3.1(iii); the deduction $\delta^{\mathsf{T}} V = 0 \Rightarrow \delta = 0$ and the general-k argument of Remark 6.3; and the census counts of §4, of which only the $n = 5$ orbit count is used by a theorem, in Theorem 3.3(ii). The complete verification costs about 45 CPU-minutes, of which a single search — the 33 592 320-map refutation of monomial equivalence at $k = 5$, $n = 6$ — is 21 minutes.

One methodological note, because it produced a wrong answer that looked like a right one. While filtering the census by distance we used the prefilter "a code of distance ≥ 3 has every coordinate of rank 2", on the ground that a rank-one coordinate i with image $\{0, a\}$ puts the weight-one label $a \cdot e_i$ into $C^{\perp}$. The prefilter is false: that label lowers the distance only when it is not itself in C, and $C = \langle a \cdot e_i \rangle \oplus$ (a perfect [[5, 1, 3]] on the other five coordinates), for instance

$$\langle \text{XIIIII, IXZZXI, IIXZZX, IXIXZZ, IZXIXZ} \rangle,$$

is full-support, self-orthogonal, jointly injective and of stabilizer distance 3 with a rank-one coordinate — checked outside the formal development, as the census itself is. With the prefilter in place the scan returned 0 distance-3 classes, and the 0 read exactly like a result; it was caught only because an earlier run of the same census had already found 1 such class at $k = 4$. The filter was removed and every count in Computation 4.2 is from the unfiltered scan, controlled first on the cell where the answer was known. A prefilter is a lemma, and an unproved one is a silent way to lose a theorem.

References

[1] Ali Assem Mahmoud, *Minimal Counterexamples of the MacWilliams Extension Theorem for Stabilizer Codes*, arXiv:2607.26214v1, submitted 28 July 2026 (quant-ph primary, math-ph cross-list). No journal version; v1 is the only version, and the internal numbers cited here were read off a compilation of its e-print.

[2] F. J. MacWilliams, *Combinatorial problems of elementary abelian groups*, Ph.D. thesis, Radcliffe College, Cambridge, MA, 1962.

[3] J. A. Wood, *Foundations of linear codes defined over finite modules: the extension theorem and the MacWilliams identities*, in: Codes over Rings (Ankara, 2008), Ser. Coding Theory Cryptol. **6**, World Scientific, 2009, pp. 124–190; `doi:10.1142/9789812837691_0004`.

[4] J. A. Wood, *Isometry groups of additive codes over finite fields*, J. Algebra Appl. **17** (2018), no. 10, 1850198; `doi:10.1142/S0219498818501980`. No e-print; the length-28 example cited above is Example 9.4 of the author's accepted manuscript.

[5] S. Dyshko, *On extendability of additive code isometries*, Adv. Math. Commun. **10** (2016), no. 1, 45–52; `doi:10.3934/amc.2016.10.45`. E-print arXiv:1406.1714v2, there titled *On extendibility of additive code isometries*; Example 3 is numbered as in v2.

[6] H. Gluesing-Luerssen and T. Pllaha, *On quantum stabilizer codes derived from local Frobenius rings*, Finite Fields Appl. **58** (2019), 145–173; arXiv:1710.09884v1; `doi:10.1016/j.ffa.2019.04.004`.

[7] T. Pllaha, *Symplectic isometries of stabilizer codes*, J. Algebra Appl. **19** (2020), no. 2, 2050021; arXiv:1807.09107v1; `doi:10.1142/S0219498820500218`.

[8] E. M. Rains, *Quantum codes of minimum distance two*, IEEE Trans. Inform. Theory **45** (1999), no. 1, 266–271; arXiv:quant-ph/9704043v1; `doi:10.1109/18.746807`.

[9] E. M. Rains, *Nonbinary quantum codes*, IEEE Trans. Inform. Theory **45** (1999), no. 6, 1827–1832; arXiv:quant-ph/9703048; `doi:10.1109/18.782103`.

[10] E. Knill and R. Laflamme, *Theory of quantum error-correcting codes*, Phys. Rev. A **55** (1997), no. 2, 900–911; arXiv:quant-ph/9604034; `doi:10.1103/PhysRevA.55.900`.

[11] D. Gottesman, *Stabilizer Codes and Quantum Error Correction*, Ph.D. thesis, California Institute of Technology, 1997; arXiv:quant-ph/9705052.

[12] A. R. Calderbank, E. M. Rains, P. W. Shor and N. J. A. Sloane, *Quantum error correction via codes over GF(4)*, IEEE Trans. Inform. Theory **44** (1998), 1369–1387; `doi:10.1109/18.681315`.

[13] P. W. Shor and R. Laflamme, *Quantum analog of the MacWilliams identities for classical coding theory*, Phys. Rev. Lett. **78** (1997), 1600–1602; `doi:10.1103/PhysRevLett.78.1600`.

[14] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: A. Platzer and G. Sutcliffe (eds.), Automated Deduction — CADE 28, Lecture Notes in Computer Science **12699**, Springer, Cham, 2021, pp. 625–635; `doi:10.1007/978-3-030-79876-5_37`.

[15] The mathlib Community, *The Lean mathematical library*, in: Certified Programs and Proofs (CPP 2020), ACM, 2020, pp. 367–381; `doi:10.1145/3372885.3373824`.