TILES WITHOUT SPECTRA IN ABELIAN 2-GROUPS: A SMALLER ELEMENTARY ABELIAN EXAMPLE, EXACT ORTHOGONALITY DEFECTS, AND THE OPTIMALITY OF A FIBRE–CLIQUE CERTIFICATE

KOREA SUPERINTELLIGENCE LABS

Abstract. Fan and Kadir have refuted the statement — formulated explicitly by Malikiosis and attributed by him to Shi — that every translational tile in a finite abelian p-group is spectral, by exhibiting three explicit tiles without spectra: a 64-point set Γ in $\mathbb{Z}_4^4 \times \mathbb{Z}_2^2$, a 512-point set R_2 in $\mathbb{F}_2^{13}$, and a 2187-point set in $\mathbb{F}_3^9$. This paper is about the first two of them.

We give a smaller counterexample. Carrying out, on the authors' own base data, the second of the two routes to smaller examples that they name in their concluding remarks — "reduce the number of layer types while forcing the weighted horizontal Fourier zero graph to have clique number below the complement size" — produces a 128-point translational tile of the elementary abelian group $\mathbb{F}_2^{11}$, of order 2048, with a 16-point tiling complement and no spectrum: the dimension drops from 13 to 11 and the ambient order from 8192 to 2048. We also show that it is the base tile, and not the choice of layer complements, that limits this method. The common Fourier zero set carrying the binary certificate is exactly the set of nonzero frequencies at which the Fourier transform of the base tile A_2 does not vanish; consequently the relevant clique number is at least 12 for *every* choice of tiling complements of A_2 and *every* layer group, so the authors' certificate is optimal for their base, with a slack of exactly 4 against the complement size 16.

Two quantitative additions accompany these. Non-spectrality asserts only that $|A|$ pairwise orthogonal characters do not exist; we compute the exact maximum in three cases — 48 of the 64 for Γ, 384 of 512 for R_2, and 96 of 128 for the new example. And for the exponent-4 example we show that its one search — that a 16-point set K does not tile $H = \mathbb{Z}_4^4 \times \mathbb{Z}_2$ — cannot be traded for a certificate of the only shape a general argument offers: a tiling complement of K is an independent set of $\mathrm{Cay}(H, (K-K) \setminus \{0\})$, so the clique–coclique bound $|C|\,|S| \leq |H|$ turns any 17-point clique of that graph into a search-free refutation of the tiling, and we prove that its clique number is exactly 16, attained by K itself. The bound then gives $|S| \leq 32$, exactly the size a tiling complement has, so it has no slack. Every theorem below is machine-checked in Lean 4; the clique numbers that are reported as measurements are marked as such and are not part of the formal development.

1. Introduction

1.1. Tiles and spectra in a finite abelian group. Let G be a finite abelian group and $\widehat{G}$ its character group. A set $A \subseteq G$ is a *translational tile* if there is $B \subseteq G$ such that every element of G has a unique representation $a + b$ with $a \in A$ and $b \in B$; one writes $A \oplus B = G$ and calls B a *tiling complement* of A. A set $A \subseteq G$ is *spectral* if there is $\Lambda \subseteq \widehat{G}$ with $|\Lambda| = |A|$ whose characters, restricted to A, form an orthogonal basis of $\mathbb{C}^A$; Λ is then a *spectrum* of A.

Fuglede's conjecture [4] in this discrete setting asks for the equivalence of the two properties. Write **T–S** for *tile* $\Rightarrow$ *spectral* and **S–T** for the converse. Both fail in general: Tao [5] disproved S–T in $\mathbb{R}^n$ for $n \geq 5$ by an example built inside a finite abelian group, and finite abelian groups have been the setting in which the two implications are separated ever since. The classical failure of T–S in a finite abelian group is the construction of Farkas, Matolcsi and Móra [6] in $\mathbb{Z}_{24}^3$. In p-groups the known failures went the other way — Ferguson and Sothanaphan [8] exhibit spectral sets that do not tile $\mathbb{F}_p^4$ for every odd p — while both implications hold in $\mathbb{Z}_p \times \mathbb{Z}_p$ [7] and in $\mathbb{F}_p^3$ for $p = 5, 7$ [9].

1.2. The p-group case. The group $\mathbb{Z}_{24}^3$ of [6], and $\mathbb{Z}_{60} \times \mathbb{Z}_{12}$, in which Zhang [11] has recently shown both directions to fail, have orders $2^9 \cdot 3^3$ and $2^4 \cdot 3^2 \cdot 5$; neither is a prime power, so neither touches the p-group case. That case had been isolated as a conjecture. Malikiosis [2, Conjecture 1.3] states it in the form

"**T-S** holds in every finite Abelian p-group,"

and attributes it there to Shi [3]; the attribution is worth preserving, since the statement is usually met under Malikiosis' name.

Fan and Kadir [1] refute it. Their Theorem 1.2 gives three explicit translational tiles without spectra: a 64-point set $\Gamma \subset \mathbb{Z}_4^4 \times \mathbb{Z}_2^2$, a 512-point set $R_2 \subset \mathbb{F}_2^{13}$, and a 2187-point set in $\mathbb{F}_3^9$, with tiling complements of sizes 16, 16 and 9. The last two are elementary abelian, so the failure is not caused by the presence of elements of order p^2; the authors record this as "the dimension upper bounds $d_2 \leq 13$ and $d_3 \leq 9$ for the first possible elementary abelian counterexamples", adding that they "do not claim that either dimension is minimal". That sentence is the only occurrence of the symbol d_p in [1], which never defines it; we read it, as its own context requires, as the least d for which $\mathbb{F}_p^d$ carries a translational tile without a spectrum, and we state our own bound below without it as well. The present paper is about the first two examples — the exponent-4 one, which is the smallest of the three, and the binary one, which is the smallest elementary abelian one they exhibit. The example in $\mathbb{F}_3^9$ is untouched here.

Two different certificates carry the two examples, and it is worth saying what each has to supply, because everything below is about the two shapes.

The exponent-4 certificate has the following shape. Inside $H = \mathbb{Z}_4^4 \times \mathbb{Z}_2$, itself identified with its dual by the pairing

$$\langle x, y \rangle = x_1y_1 + x_2y_2 + x_3y_3 + x_4y_4 + 2x_5y_5 \pmod 4, \qquad \chi_y(x) = i^{\langle x,y \rangle}, \tag{1}$$

the authors exhibit four sets K, E, T_0, T_1 — $|K| = |E| = 16$, $|T_0| = |T_1| = 32$ — and prove about them the four assertions of their Proposition 5.1: (1) E is a spectrum of K; (2) $E \oplus T_0 = H$ and $E \oplus T_1 = H$; (3) K does not tile H; (4) for every nonzero $d \in K - K$ one has $|\widehat{1_{T_0}}(d)|^2 \neq |\widehat{1_{T_1}}(d)|^2$. Their Lemma 3.1 then assembles $\Gamma = (T_0 \times \{0\}) \cup (T_1 \times \{1\}) \subset H \times \mathbb{Z}_2 \cong \mathbb{Z}_4^4 \times \mathbb{Z}_2^2$. Items (1), (2) and (4) are finite verifications on 16×16, 16×32 and 152 items respectively. Item (3) is not: it is a statement that *no* subset of H of any shape completes K to a tiling, and it is proved by exhausting an exact-cover tree, for which the authors report 42 nodes.

The binary certificate has a different shape, and its one search is a search a reader can run in a second. Inside $H_2 = \mathbb{F}_2^8$ the authors exhibit a 16-point set A_2 and seventeen 16-point sets $C_0, \ldots, C_{16}$, each of them a tiling complement of A_2; weights $(a_0, \ldots, a_{16})$ summing to 32 distribute the C_j over the layers of $L_2 = \mathbb{F}_2^5$, and $R_2 = \bigcup_j (C_j \times Y_j) \subset \mathbb{F}_2^8 \times \mathbb{F}_2^5 = \mathbb{F}_2^{13}$. Their Lemma 4.1 — the *fibre–clique* obstruction — then says: if the Cayley graph on $\mathbb{F}_2^8$ whose connection set D_2 is the zero set of the weighted transform $W_2 = \sum_j a_j \widehat{1_{C_j}}$ has no clique of size $q = |C_j| = 16$, then R_2 tiles $\mathbb{F}_2^{13}$ and is not spectral. Their Proposition 6.1 supplies the four finite facts that instantiate it: the seventeen tilings; $|D_2| = 108$, with D_2 the exact zero set of W_2 away from the origin; the clique number $\omega(\mathrm{Cay}(\mathbb{F}_2^8, D_2)) = 12$; and a magnitude bound off D_2. Only the third is a search, and it is a search on 256 vertices.

1.3. **What this paper adds.** Five things, of which the first two are the point.

(a) A smaller elementary abelian counterexample. [1] closes by naming two routes to smaller examples: reduce the dimension of the base group, or "reduce the number of layer types while forcing the weighted horizontal Fourier zero graph to have clique number below the complement size". We carry out the second on the authors' own base data. The observation that makes it a small search is that the inequality the fibre–clique lemma needs, $\omega(\mathrm{Cay}(X, D)) < q$, does not mention the layer group at all, so shrinking the layer group does not weaken it; and the weighted transform depends only on the *multiset* of complements assigned to the layers. With A_2 and the seventeen printed complements kept fixed and the layer group taken to be $\mathbb{F}_2^3$, exactly three of the 735,471 multisets of size 8 work. One of them gives *a* 128*-point translational tile* R_S *of* $\mathbb{F}_2^{11}$*, of order* 2048*, with tiling complement* $A_2\times\{0\}$ *of size* 16 *and no spectrum* (Theorem 8.1).

This lowers the authors' upper bound on the least dimension of an elementary abelian binary counterexample from 13 to 11, and the ambient order from 8192 to 2048. It uses no idea that is not in [1]; what it adds is the observation about $|Y|$, and the search.

(b) The base tile is the bottleneck, and the authors' certificate is optimal for it. D_2 is defined in [1] as the intersection $\bigcap_{j=0}^{16} Z(C_j)$ of seventeen zero sets. In fact the seventeen complements contribute nothing to it:

$D_2 = \{v \neq 0 : \widehat{1_{A_2}}(v) \neq 0\}$*, and consequently* $\omega(\mathrm{Cay}(\mathbb{F}_2^8, D)) \geq 12$ *for every layer group and every assignment of tiling complements of* A_2 *to its layers* (Theorem 9.2).

The second half follows from the first through the tiling identity $\widehat{1_A}\cdot\widehat{1_C}\equiv 0$ off the origin: whatever complements are chosen, the weighted transform is forced to vanish on all of D_2. Since the authors' own weights achieve exactly 12, their certificate cannot be improved with this base: against $q = 16$ the slack is exactly 4. The contrast with the exponent-4 example is sharp — there the corresponding clique number is $16 = q$ on the nose and the clique route is closed off entirely, which is why two different lemmas are needed.

(c) How far the tiles are from spectral. Non-spectrality says that no $|A|$ characters of the ambient group are pairwise orthogonal on A. It says nothing about how many there are, and [1] computes no such quantity. We compute the exact maximum three times:

Γ *carries exactly* 48 *pairwise orthogonal characters of the* 64 *a spectrum would need;* R_2 *exactly* 384 *of* 512*; and* R_S *exactly* 96 *of* 128 (Theorems 5.1, 7.7 and 8.1, with Remark 5.2).

For R_2 and R_S both halves are machine-checked; for Γ the lower half is, and the matching upper bound 48 is a computation reported outside the formal development (Remark 5.2). The three ratios are all 3/4, which is an observation on three data points and nothing more (Remark 8.3).

(d) The search of the exponent-4 certificate cannot be replaced by a certificate. There is exactly one general tool that could refute the tiling of item (3) without a search. Put $D_K = (K - K)\setminus\{0\}$ and consider the Cayley graph $\mathrm{Cay}(H, D_K)$. A tiling complement S of K is an independent set of that graph, while any clique C of it satisfies $|C|\,|S| \leq |H| = 512$ by the clique–coclique inequality. Since a tiling complement of K must have $512/16 = 32$ elements, a clique of 17 points would contradict the existence of one outright, and would do so in a form a reader can check by hand. We prove that no such clique exists.

The clique number of $\mathrm{Cay}(H, D_K)$ *is exactly* 16*, attained by* K *itself* (Theorems 4.3 and 4.4).

So the clique–coclique bound gives $|S| \leq 32$ and nothing better, which is precisely the size a tiling complement has: the bound has no slack, and the exhaustive search of item (3) is not an artefact of the authors' method. The upper bound is itself established by an exhaustive enumeration, of 126,295 nodes; but it is an enumeration of cliques of a 152-vertex induced subgraph, and what it buys is the statement that the cheap route is closed.

(e) The certificates, reproduced, and one of them shortened. Every number that [1] prints for the two examples treated here is recomputed from the coordinate data of its appendices, and every one of them that is an assertion about that data is then machine-checked (Sections 3, 7 and 11). The two numbers it reports that depend on an algorithm rather than on the data are treated apart: the exact-cover node count of the exponent-4 verifier is reproduced but not verified (Remark 11.4), and the node count reported for the colouring branch-and-bound of the binary example is neither reproduced nor verified, the search replayed here having a different shape (Remark 11.5). In the course of that, a shorter certificate for the exponent-4 example falls out. Of the four items of Proposition 5.1, only item (3) is used below. Items (1) and (2) enter the argument of [1] to produce the tiling of Γ and the count $|\Lambda| = 2|H|/|K|$; here the tiling is checked directly and the count is read off $|\Gamma| = 64$. Item

(4) enters only through $\widehat{1_\Gamma}(d,0) = \widehat{1_{T_0}}(d) + \widehat{1_{T_1}}(d)$, to force $\widehat{1_\Gamma}(d,0) \neq 0$ for nonzero $d \in K - K$; here that conclusion is checked directly on Γ. What remains is the triple (Γ, E, K) with two finite checks and one exact-cover search, and no property of the two-layer decomposition of Γ is needed. This is a simplification of the authors' own proof and is presented as such: it uses no idea that is not in [1]. Items (1), (2) and (4) are reproduced anyway in Section 11, because they are the certificate the authors print.

1.4. **What is not claimed.** The disproof of the p-group conjecture is that of [1], and nothing in Sections 3 and 7 is offered as new mathematics: the sets are theirs, the two obstruction lemmas are theirs, and what is added is that the whole chain, including both searches, is checked by a proof assistant. The smaller example of Section 8 is a smaller instance of their construction, found by carrying out a route they name themselves; it is new, and it is not a new mechanism. Nothing is claimed about their third counterexample, in $\mathbb{F}_3^9$; its clique step is larger and is untouched here. We do not determine the least dimension of an elementary abelian 2-group carrying a tile without a spectrum: 11 is an upper bound for it, not a value, and Remark 8.4 says exactly which parts of the search behind Theorem 8.1 are complete and which are not. Nor do we determine the least order of an abelian p-group carrying a tile without a spectrum — [1] explicitly leaves this open, and 1024 is only the order of the smallest example now known. Finally, Theorems 4.3 and 4.4 say that the clique–coclique bound cannot refute the tiling of K; they do not say that no short human proof of item (3) exists, only that this particular general mechanism yields none.

2. Preliminaries

2.1. **The groups, the pairing, the characters.** Throughout, $H = \mathbb{Z}_4^4 \times \mathbb{Z}_2$, of order 512, and

$$G = H \times \mathbb{Z}_2 \;\cong\; \mathbb{Z}_4^4 \times \mathbb{Z}_2^2,$$

of order $1024 = 2^{10}$; the isomorphism is the reassociation of the product, and it is used without comment. We write elements of H in the notation $(abcd; e)$ of [1], so that $(3202; 0)$ means $(3, 2, 0, 2; 0)$.

The pairing on H is (1); on G it is extended by the term coming from the second $\mathbb{Z}_2$ factor,

$$\langle x, y\rangle_G = \langle x', y'\rangle + 2x_6y_6 \pmod 4, \qquad x = (x', x_6),\ y = (y', y_6) \in H \times \mathbb{Z}_2, \tag{2}$$

which is the self-dual identification of $\widehat{H} \times \widehat{\mathbb{Z}_2}$ with $H \times \mathbb{Z}_2$ used in [1] when a spectrum of Γ is written as a subset of $H \times \mathbb{Z}_2$. The pairing (2) is nondegenerate (Proposition 2.4), so

$$\chi_y(x) = i^{\langle x,y\rangle_G} \qquad (y \in G)$$

runs over the 1024 distinct characters of G as y runs over G, and a subset $\Lambda \subseteq G$ names $|\Lambda|$ distinct characters. This is the only role nondegeneracy plays, and it is not cosmetic: without it a 64-element subset of G would not be a set of 64 characters.

2.2. **Fourier coefficients as Gaussian integers.** For $A \subseteq G$ and $w \in G$ put $\widehat{1_A}(w) = \sum_{a\in A} \overline{\chi_w(a)}$, and let

$$c_r(A, w) = \#\{a \in A : \langle a, w\rangle_G = r\} \qquad (r \in \mathbb{Z}_4).$$

Then $\sum_{a\in A} \chi_w(a) = (c_0 - c_2) + (c_1 - c_3)\, i$ is a Gaussian integer, and $\widehat{1_A}(w)$ is its complex conjugate. In particular:

Lemma 2.1 (the exact vanishing test). *$\widehat{1_A}(w) = 0$ if and only if $c_0(A, w) = c_2(A, w)$ and $c_1(A, w) = c_3(A, w)$.*

Every finite check in this paper is stated through the integer condition on the right, and Lemma 2.1 is what makes it a statement about the complex Fourier coefficient rather than about a surrogate for it. We write $|\widehat{1_A}(w)|^2 = (c_0 - c_2)^2 + (c_1 - c_3)^2$, again an exact integer; this is the form in which [1] says its own verifier compares magnitudes. The same definitions are used on H, with $\langle\cdot,\cdot\rangle$ in place of $\langle\cdot,\cdot\rangle_G$.

2.3. **The two criteria.** Two standard facts are used as stated in [1].

Lemma 2.2 (difference test)**.** *Let X be a finite abelian group and $A, B \subseteq X$ with $|A|\,|B| = |X|$. If $(A - A) \cap (B - B) = \{0\}$, then $A \oplus B = X$.*

Lemma 2.3 (spectral test)**.** *$\Lambda \subseteq G$ is a spectrum of $A \subseteq G$ if and only if $|\Lambda| = |A|$ and $\widehat{1_A}(u - v) = 0$ for all distinct $u, v \in \Lambda$.*

Lemma 2.3 is the reason spectrality can be treated as a purely combinatorial condition on differences; it is the criterion [1] uses, and orthogonality of the restricted characters is what "orthogonal basis of $\mathbb{C}^A$" amounts to once the cardinality is fixed. We also use repeatedly that if $|A|\,|B| = |X|$ and every element of X is *some* sum $a + b$, then the representation is automatically unique, so $A \oplus B = X$. Lemma 2.3 is used below in each of the three ambient groups that occur — $\mathbb{Z}_4^4 \times \mathbb{Z}_2^2$, $\mathbb{F}_2^{13}$ and $\mathbb{F}_2^{11}$ — with the self-dual identification named there, and in the formal development it is proved separately in each, from the corresponding exact vanishing test, since the vanishing test is the part that differs. The counting fact just stated, and Lemma 2.2, are proved once for an arbitrary finite abelian group and reused; Lemma 2.2 is used only for the exponent-4 example.

2.4. **The data.** The four sets $K, E, T_0, T_1 \subseteq H$ are those of Appendix A of [1] and are reproduced in Appendix A below; they were extracted from the source's own LaTeX by a script matching its coordinate macros, so no hand transcription occurs anywhere in the chain. Set

$$\Gamma = (T_0 \times \{0\}) \cup (T_1 \times \{1\}) \subseteq G, \qquad E_0 = E \times \{0\}, \qquad K_0 = K \times \{0\}.$$

Proposition 2.4 (the data)**.** *$|K| = |E| = 16$, $|T_0| = |T_1| = 32$, $|\Gamma| = 64$, $|E_0| = |K_0| = 16$, and the pairing* (2) *is nondegenerate: for every $y \neq 0$ in G there is $x \in G$ with $\langle x, y\rangle_G \neq 0$.*

Proof sketch. The cardinalities are decided directly from the lists, so in particular the lists have no repetitions. Nondegeneracy is an exhaustive check over the 1024×1024 pairs. □

3. The exponent-4 counterexample, and a certificate with three ingredients

This section restates the counterexample of [1] for the group $\mathbb{Z}_4^4 \times \mathbb{Z}_2^2$ and records that each step has been machine-checked. The mathematics is that of [1]; the arrangement isolates what the conclusion actually consumes.

Theorem 3.1 (the tiling; [1], Theorem 1.2(1))**.** *$\Gamma \oplus (E \times \{0\}) = G$. In particular Γ is a translational tile of $\mathbb{Z}_4^4 \times \mathbb{Z}_2^2$.*

Proof sketch. $|\Gamma|\,|E_0| = 64 \cdot 16 = 1024 = |G|$, and every $g \in G$ is checked to be some sum $a + b$ with $a \in \Gamma$, $b \in E_0$; that check is exhaustive over the 1024 elements of G. Uniqueness is then a counting consequence and is not itself checked. The authors obtain the same tiling indirectly, from $E \oplus T_0 = H$ and $E \oplus T_1 = H$ (Proposition 11.2). □

Theorem 3.2 ([1], Proposition 5.1(3))**.** *K has no tiling complement in $H = \mathbb{Z}_4^4 \times \mathbb{Z}_2$.*

Proof sketch. This is the one step that is a search, and it is worth saying exactly what was run. A tiling complement S would have $|S| = 512/16 = 32$ elements, and by translating S we may assume $0 \in S$, which is the normalisation [1] describes. So the question is an exact cover: starting from the single translate $K + 0$, can 31 further translates $K + t$, pairwise disjoint, cover the remaining 496 elements of H? The state of the search is a 512-bit integer whose bit at the mixed-radix code of $x \in H$ records that x is covered; at each node the lowest uncovered code is resolved by trying the 16 translates that could cover it. The procedure returns *false*.

Two things are proved rather than assumed. First, that the procedure is complete: *if* a tiling complement exists, the search returns *true*. The proof uses only that the branching function returns *some* uncovered index, so the tree replayed by the machine need not be the one [1] explored. Second, that the bit-mask state really represents the set of chosen translates; this invariant is the only place where the machine word and the subset of H meet. The exhaustive part of the argument is therefore exactly the Boolean verdict of one finite search, and its size — 24,478 nodes with the branching rule just described, measured on an identical implementation outside the formal development — is a cost, not a claim. Section 11 compares it with the 42 nodes [1] prints. □

Lemma 3.3. *For all distinct $u, v \in K \times \{0\}$ one has $\widehat{1_\Gamma}(u - v) \neq 0$; equivalently, $\widehat{1_\Gamma}(d, 0) \neq 0$ for every nonzero $d \in K - K$.*

Proof sketch. An exhaustive check of the integer test of Lemma 2.1 over the 240 ordered pairs of distinct elements of K. In [1] the same conclusion is reached from Proposition 5.1(4) through $\widehat{1_\Gamma}(d, 0) = \widehat{1_{T_0}}(d) + \widehat{1_{T_1}}(d)$; evaluating one Fourier transform on Γ instead of two on T_0, T_1 is the same fact obtained more directly, and it removes the need for the two-layer structure. □

Theorem 3.4 (non-spectrality; [1], Theorem 1.2(1))**.** *Γ has no spectrum: no 64 characters of $\mathbb{Z}_4^4 \times \mathbb{Z}_2^2$ are pairwise orthogonal on Γ.*

Proof sketch. Suppose $\Lambda \subseteq G$ is a spectrum, so $|\Lambda| = |\Gamma| = 64$ and, by Lemma 2.3, every nonzero difference of Λ is a zero of $\widehat{1_\Gamma}$. By Lemma 3.3 no nonzero difference of $K_0 = K \times \{0\}$ is such a zero, so $(\Lambda - \Lambda) \cap (K_0 - K_0) = \{0\}$; since $64 \cdot 16 = 1024 = |G|$, Lemma 2.2 gives $\Lambda \oplus K_0 = G$.

Now restrict to the two layers. For $l \in \mathbb{Z}_2$ let $\Lambda_l = \{y \in H : (y, l) \in \Lambda\}$. Every $(h, l) \in G$ is uniquely $\lambda + \kappa$ with $\lambda \in \Lambda$ and $\kappa \in K_0$, and κ has second coordinate 0; hence $\lambda \in \Lambda_l \times \{l\}$ and $K + \Lambda_l = H$ for each l. Therefore $|K|\,|\Lambda_l| \geq |H|$, i.e. $|\Lambda_l| \geq 32$, for both l; as $|\Lambda_0| + |\Lambda_1| = 64$, both are exactly 32. Then $|K|\,|\Lambda_0| = 512 = |H|$ together with $K + \Lambda_0 = H$ forces $K \oplus \Lambda_0 = H$, contradicting Theorem 3.2.

This is a repackaging of Lemma 3.1 of [1], whose hypotheses are all four items of its Proposition 5.1; here only item (3) is used, via Theorem 3.2. □

Corollary 3.5 ([1])**.** *There is a subset of the abelian 2-group $\mathbb{Z}_4^4 \times \mathbb{Z}_2^2$, of order $1024 = 2^{10}$, that is a translational tile and is not spectral. Hence T–S fails in a finite abelian p-group, and the conjecture of* [2, Conjecture 1.3]*, attributed there to* [3]*, is false.*

Proof. Theorems 3.1 and 3.4, with $A = \Gamma$. That the ambient group is a 2-group is a property of the group itself, not a hypothesis of the statement: the formal statement quantifies over subsets of this one concrete group of order 2^{10}. □

4. The clique number of $\mathrm{Cay}(H, D_K)$

Theorem 3.2 is the only step of Section 3 that is not a bounded verification of a printed object, and it is natural to ask whether it can be replaced by one. This section answers no, for the one mechanism a general argument supplies.

Write $D_K = (K - K) \setminus \{0\}$, a symmetric subset of H with $|D_K| = 152$ (Proposition 11.3), and let $\mathrm{Cay}(H, D_K)$ be the Cayley graph on vertex set H in which $x \sim y$ iff $x - y \in D_K$. A *clique* is a subset whose distinct elements have all their differences in D_K; an *independent set* is one no two of whose distinct elements differ by an element of D_K.

Lemma 4.1 (clique–coclique bound)**.** *Let X be a finite abelian group, $D \subseteq X$, and let $C, S \subseteq X$ be such that $c - c' \in D$ for all distinct $c, c' \in C$ and $s - s' \notin D$ for all distinct $s, s' \in S$. Then $|C|\,|S| \leq |X|$.*

Proof. The map $(c, s) \mapsto c - s$ is injective on $C \times S$: if $c - s = c' - s'$ with $c \neq c'$ then $s \neq s'$ and $c - c' = s - s'$ lies in D and not in D; and if $c = c'$ then $s = s'$. Hence $|C|\,|S| = |C \times S| \leq |X|$. □

For vertex-transitive graphs this is the classical clique–coclique inequality (a standard fact of algebraic graph theory; see e.g. [12]); on an abelian Cayley graph it needs no machinery, and the two lines above are the whole proof.

Proposition 4.2. *K is a clique of $\mathrm{Cay}(H, D_K)$ with 16 vertices, and every tiling complement of K is an independent set of $\mathrm{Cay}(H, D_K)$. Consequently, if K had a tiling complement S, then $|S| = 32$ and every clique C of $\mathrm{Cay}(H, D_K)$ would satisfy $|C| \leq 16$.*

Proof. That distinct $k, k' \in K$ have $k - k' \in D_K$ is immediate from the definition of D_K and is checked exhaustively over the 240 ordered pairs. If S is a tiling complement and $s, s' \in S$ are distinct with $s - s' = k - k' \in D_K$ for some $k \neq k'$ in K, then $k + s' = k' + s$ are two representations of one element, so uniqueness forces $(k, s') = (k', s)$, contradicting $k \neq k'$. The last clause is Lemma 4.1 with $X = H$ and $D = D_K$, together with $|K|\,|S| = |H|$. □

The conditional in the last sentence has a false hypothesis once Theorem 3.2 is known; what carries content is its contrapositive, together with the unconditional bound $|S| \leq 32$ of Corollary 4.5, which needs only Lemma 4.1 and the 16-clique K. Read that way, Proposition 4.2 is a proof strategy: exhibit 17 points of H whose pairwise differences all lie in $K - K$, and Theorem 3.2 follows with no search — a certificate of 17 coordinates that a reader can verify by hand in a few minutes. The next theorem is that no such certificate exists.

Theorem 4.3. *No 17-point subset of $H = \mathbb{Z}_4^4 \times \mathbb{Z}_2$ has all of its pairwise differences in $K - K$. Equivalently, every clique of $\mathrm{Cay}(H, D_K)$ has at most 16 vertices.*

Proof sketch. Cliques of a Cayley graph are translation invariant, so a clique of 17 points may be translated to contain 0; its other 16 points then lie in D_K and form a 16-clique inside the 152-vertex induced subgraph on D_K. It therefore suffices to show that D_K contains no 16-clique, and *this is the claim that rests on exhaustive computation.*

The finite search is the plain increasing-index clique enumeration over the 152 elements of D_K in a fixed order: at a node one either takes the current vertex and recurses on the sublist of its later neighbours, or skips it, with the single prune "a list shorter than the remaining target contains no clique of that size". Adjacency is a single bit of a 512-bit word indexed by the mixed-radix code of the difference. The enumeration returns *false*, and the completeness of the enumeration — if a 16-clique existed among the listed vertices, it would return *true* — is proved, using no property of the graph. Measured on an identical implementation outside the formal development, the tree has 126,295 nodes and makes 921,430 adjacency tests. A separate branch-and-bound with a greedy-colouring bound, also outside the formal development, confirms the same conclusion: the subgraph induced on D_K has clique number 15, so $\mathrm{Cay}(H, D_K)$ has clique number 16, in 227 nodes. □

Theorem 4.4. *The clique number of $\mathrm{Cay}(H, D_K)$ is exactly 16, attained by K.*

Proof. The lower bound is Proposition 4.2 and the upper bound is Theorem 4.3. □

Corollary 4.5. *Every independent set S of $\mathrm{Cay}(H, D_K)$ satisfies $|S| \leq 32$, and no bound of clique–coclique type gives less. Since a tiling complement of K would be an independent set with exactly 32 elements, no clique of $\mathrm{Cay}(H, D_K)$ contradicts the existence of one: Theorem 3.2 does not follow from Lemma 4.1.*

Proof. Lemma 4.1 applied to the clique K gives $16\,|S| \leq 512$; no larger clique is available, by Theorem 4.4. A tiling complement is independent by Proposition 4.2 and has $|H|/|K| = 32$ elements. □

Remark 4.6 (the same tightness one level up, in part measured). The same phenomenon occurs for Γ itself, and explains why the layer argument in the proof of Theorem 3.4 — or the two-layer Lemma 3.1 of [1] — is needed rather than a single-set certificate. Let $Z(\Gamma) = \{w \neq 0 : \widehat{1_\Gamma}(w) = 0\}$ and $N(\Gamma) = \{w \neq 0 : \widehat{1_\Gamma}(w) \neq 0\}$, so that $|Z(\Gamma)| = 437$ and $|N(\Gamma)| = 586$. By Lemma 2.3 a spectrum of Γ is a 64-clique of $\mathrm{Cay}(G, Z(\Gamma))$, and an independent set of that graph is exactly a clique of $\mathrm{Cay}(G, N(\Gamma))$. Lemma 3.3 says precisely that $K \times \{0\}$ is a 16-clique of $\mathrm{Cay}(G, N(\Gamma))$, so Lemma 4.1, applied with that independent set, bounds a clique of $\mathrm{Cay}(G, Z(\Gamma))$ by $1024/16 = 64$ — exactly the size a spectrum would need, so no contradiction is produced. *Computed outside the formal development*, by a branch-and-bound with a greedy-colouring bound: the clique number of $\mathrm{Cay}(G, N(\Gamma))$ is exactly 16, the neighbourhood of 0 requiring 9,098,430 nodes; so no independent set of $\mathrm{Cay}(G, Z(\Gamma))$ does better than $K \times \{0\}$, and the bound cannot be improved by a different choice. Of this remark, only the lower bound 16 belongs to the formal development.

5. How far Γ is from spectral

Theorem 5.1. *Let $\Omega \subset H$ be the* 24*-point set listed in Appendix A and put $\Lambda_{48} = \Omega \times \mathbb{Z}_2 \subseteq G$. Then $|\Lambda_{48}| = 48$ and the characters χ_u, $u \in \Lambda_{48}$, are pairwise orthogonal on Γ:*

$$\sum_{a \in \Gamma} \chi_u(a)\,\overline{\chi_v(a)} = 0 \qquad \text{for all distinct } u, v \in \Lambda_{48}.$$

Since Γ has no spectrum (Theorem 3.4), which would consist of 64 *pairwise orthogonal characters, Γ carries at least* 48 *pairwise orthogonal characters and no* 64.

Proof sketch. $|\Lambda_{48}| = 48$ is decided from the list. The $\binom{48}{2} = 1128$ inner products are evaluated as complex sums: each is rewritten as the single Fourier coefficient $\sum_{a \in \Gamma} \chi_{u-v}(a)$, and its vanishing is decided by the exact integer test of Lemma 2.1. *That check —* 1128 *vanishing tests — is the exhaustive part*; everything else is the identity $\chi_u \overline{\chi_v} = \chi_{u-v}$ and Lemma 2.1. □

Remark 5.2 (48 is the maximum; measured). Outside the formal development we have computed the exact maximum. A family of characters pairwise orthogonal on Γ is a clique of $\mathrm{Cay}(G, Z(\Gamma))$, a 1024-vertex graph with $|Z(\Gamma)| = 437$; by vertex-transitivity the question is the clique number of the neighbourhood of 0, which a branch-and-bound with a greedy-colouring bound determines to be 47 in 470,426 nodes. Hence the maximum number of characters pairwise orthogonal on Γ is exactly 48, and Λ_{48} attains it. So Γ misses spectrality by exactly sixteen characters. This upper bound is a computation reported here, not part of the formal development: the enumeration of Theorem 4.3, which carries a completeness proof but no colouring prune, passes 3×10^7 nodes on this instance.

6. The fibre–clique obstruction

The remaining sections concern the elementary abelian side, where the mechanism is not the two-layer lemma of [1] but its fibre–clique lemma. We state that lemma in the two halves in which it is used below, both for arbitrary finite abelian groups, because the second example uses them at a different layer group from the first.

Throughout this section X (the *base*) and Y (the *layer group*) are finite abelian groups, $A \subseteq X$, and $C_y \subseteq X$ is given for each $y \in Y$. Put

$$R = \bigcup_{y \in Y} (C_y \times \{y\}) \subseteq X \times Y.$$

Lemma 6.1 ([1], Lemma 4.1, tiling half)**.** *If $A \oplus C_y = X$ for every $y \in Y$, then $(A \times \{0\}) \oplus R = X \times Y$. In particular R is a translational tile of $X \times Y$ with tiling complement $A \times \{0\}$, and $|R| = \sum_{y \in Y} |C_y|$.*

Proof. This is the authors' "layer by layer" step; it is written out only to record what it costs. Given $(x, y) \in X \times Y$, the tiling $A \oplus C_y = X$ writes $x = a + c$ uniquely with $a \in A$ and $c \in C_y$, and $(x, y) = (a, 0) + (c, y)$ with $(c, y) \in R$. Conversely, in a representation $(x, y) = (a, 0) + r$ with $a \in A$ and $r \in R$ the second coordinate forces $r = (c, y)$ with $c \in C_y$, and the first forces $a + c = x$; so it is the representation just produced. No Fourier analysis enters, and $X \times Y$ is never enumerated: the cost is the $|Y|$ tilings of the base. □

Lemma 6.2 ([1], Lemma 4.1, clique half, quantified)**.** *Let $D \subseteq X$ contain every nonzero $v \in X$ for which $(v, 0)$ is a zero of $\widehat{1_R}$, and let w be a bound on the sizes of the cliques of* $\mathrm{Cay}(X, D)$. *Then every $\Lambda \subseteq X \times Y$ whose nonzero differences are all zeros of $\widehat{1_R}$ satisfies*

$$|\Lambda| \;\le\; |Y|\, w.$$

In particular, if $|Y|\, w < |R|$ then R has no spectrum.

Proof. Partition Λ by its second coordinate. Two distinct elements of one fibre differ by $(v, 0)$ with $v \neq 0$; that difference is a zero of $\widehat{1_R}$ by hypothesis, so $v \in D$. Hence the first coordinates of a fibre are pairwise distinct and form a clique of $\mathrm{Cay}(X, D)$, which has at most w vertices; summing over the

$|Y|$ fibres gives the bound. The last sentence is the spectral test: a spectrum of R is such a Λ with $|\Lambda| = |R|$.

The argument is that of [1], with one observation added: it never uses the size of Λ, so it bounds *every* pairwise orthogonal family and not only one of size $|R|$. [1] states the conclusion in the unquantified form: if the Cayley graph has no clique of size $q = |C_y|$, then R tiles but is not spectral. The quantitative form is what Theorems 7.7 and 8.1 use. In the formal development the statement is purely combinatorial: "is a zero of $\widehat{1_R}$" is an arbitrary property of $X \times Y$ there, so no character theory is assumed for it. $\square$

7. The binary example: $R_2 \subset \mathbb{F}_2^{13}$

7.1. **The groups, the characters, the data.** Write $H_2 = \mathbb{F}_2^8$, $L_2 = \mathbb{F}_2^5$ and $G_2 = H_2 \times L_2 = \mathbb{F}_2^{13}$, of orders 256, 32 and 8192; the identification of the product with the flat 13-fold product is the reassociation and is used without comment. Following [1] an element of $\mathbb{F}_2^d$ is named by its integer code $x = \sum_{i<d} x_i 2^i$, addition of codes being coordinatewise modulo two. Each of the groups is identified with its dual by the standard dot product, and the characters are real:

$$\chi_w(x) = (-1)^{\langle x,w\rangle}, \qquad \widehat{1_A}(w) = \sum_{a\in A}(-1)^{\langle a,w\rangle} = c_0(A,w) - c_1(A,w),$$

where $c_r(A,w) = \#\{a \in A : \langle a,w\rangle = r\}$ for $r \in \mathbb{F}_2$. So a Fourier coefficient here is an integer, and it vanishes exactly when the two fibre counts of $\langle\cdot,w\rangle$ on A agree: this is the exact test that plays for the binary examples the role Lemma 2.1 plays for Γ, and again every finite check below is stated through it while the conclusions are about the complex characters. The dot product on $\mathbb{F}_2^{13}$ is nondegenerate (Proposition 7.1), so a subset of $\mathbb{F}_2^{13}$ names that many distinct characters.

Appendix B below reproduces, from Appendix B of [1], a 16-point set $A_2 \subset H_2$, seventeen 16-point sets $C_0, \ldots, C_{16} \subseteq H_2$, and the weights

$$(a_0, \ldots, a_{16}) = (1,1,2,1,1,1,3,1,1,2,3,3,1,5,3,2,1), \qquad \textstyle\sum_j a_j = 32. \tag{3}$$

The layer group L_2 is partitioned, in increasing code order, into consecutive blocks Y_j with $|Y_j| = a_j$, and

$$R_2 = \bigcup_{j=0}^{16}(C_j \times Y_j) \subseteq \mathbb{F}_2^8 \times \mathbb{F}_2^5 = \mathbb{F}_2^{13}.$$

Finally, with $Z(S) = \{v \neq 0 : \widehat{1_S}(v) = 0\}$,

$$W_2(v) = \sum_{j=0}^{16} a_j\,\widehat{1_{C_j}}(v), \qquad D_2 = \bigcap_{j=0}^{16} Z(C_j).$$

These data were extracted from the LaTeX source of [1] by a script reading the rows of its appendix table, so no hand transcription occurs in this chain either.

Proposition 7.1 (the data)**.** *$\mathbb{F}_2^{13}$ has order 8192 and satisfies $x + x = 0$ for every x; $|A_2| = 16$; $|C_j| = 16$ for each of the seventeen indices j; $|R_2| = 512$ and $|A_2 \times \{0\}| = 16$; and the dot product on $\mathbb{F}_2^{13}$ is nondegenerate.*

Proof sketch. The cardinalities are decided from the lists, so in particular no list has a repetition — except $|R_2| = 512$, which is obtained as $\sum_{y\in L_2}|C_y| = 32 \cdot 16$ from Lemma 6.1 rather than by enumerating R_2. Nondegeneracy is an exhaustive check of the thirteen standard basis vectors against the 8192 frequencies. $\square$

7.2. **The authors' certificate, machine-checked.**

Proposition 7.2 ([1], Proposition 6.1(1))**.** *$A_2 \oplus C_j = H_2 = \mathbb{F}_2^8$ for every j with $0 \le j \le 16$.*

Proof sketch. $|A_2|\,|C_j| = 16 \cdot 16 = 256 = |H_2|$, and each of the seventeen pairs is checked to cover H_2; *that check — $17 \cdot 256$ elements — is the exhaustive part.* Uniqueness is then a counting consequence and is not itself checked. $\square$

Proposition 7.3 ([1], Proposition 6.1(2), and its identity at horizontal frequencies)**.** $|D_2| = 108$*; for every nonzero* $v \in H_2$ *one has* $W_2(v) = 0$ *if and only if* $v \in D_2$*; and* $\widehat{1_{R_2}}(v, 0) = W_2(v)$ *for every* $v \in H_2$*. Consequently* D_2 *is exactly the set of nonzero* $v \in H_2$ *for which* $(v, 0)$ *is a zero of* $\widehat{1_{R_2}}$*.*

Proof sketch. Four exhaustive checks over the 256 frequencies of H_2, all in exact integer arithmetic: that the listed 108-element set is $\bigcap_j Z(C_j)$; that it has 108 elements, the authors' printed `common_fourier_zeros=108`; that W_2 vanishes on it and nowhere else off the origin; and the identity $\widehat{1_{R_2}}(v, 0) = W_2(v)$, which is the authors' own equation for horizontal frequencies. The last clause combines the second and the fourth, and it is the only property of D_2 that the non-spectrality argument consumes. □

Proposition 7.4 ([1], Proposition 6.1(4))**.** $|W_2(v)| \geq 4$ *for every* $v \notin D_2 \cup \{0\}$*, and the value* 4 *is attained.*

Proof sketch. Two exhaustive checks over the 256 frequencies, in exact integers; together they are the authors' printed `weighted_noncancellation=1 minimum_abs_weighted_fourier=4`. This proposition is reproduction: nothing below uses it. □

Theorem 7.5 ([1], Proposition 6.1(3))**.** *The clique number of* $\mathrm{Cay}(\mathbb{F}_2^8, D_2)$ *is exactly* 12*, attained by the* 12*-point set* Θ *of Appendix B.*

Proof sketch. The lower bound is a check of the $\binom{12}{2}$ differences of Θ against D_2. For the upper bound, cliques of a Cayley graph are translation invariant, so a clique of 13 points may be translated to contain 0; its other 12 points then lie in D_2 and form a 12-clique of the induced subgraph on D_2. *So the claim that rests on exhaustive computation is that the* 108*-point set* D_2 *contains no* 12*-clique.* The finite search is the plain increasing-index clique enumeration used in Theorem 4.3, run over the 108 elements of D_2 in a fixed order, with the single prune "a list shorter than the remaining target contains no clique of that size" and with adjacency a single bit of a 256-bit word. It returns *false*, and the completeness of the enumeration — if a 12-clique existed among the listed vertices, it would return *true* — is proved, using no property of the graph. *No colouring bound is used, and at this size none is needed*: measured on an identical implementation outside the formal development, the tree has 26,683 nodes and makes 171,696 adjacency tests. That the enumeration is not a procedure which always answers *false* is checked separately: at target 11 it answers *true*, in 50 nodes (Proposition 10.2(v)). For the same conclusion [1] reports "an exact coloring branch-and-bound algorithm on 256 vertices" that "exhibits a 12-clique and exhaustively rules out a 13-clique in 425 search nodes", and an independent binary MILP formulation returning the same optimum. □

Theorem 7.6 ([1], Theorem 1.2(2))**.** $|R_2| = 512$*,* $R_2 \oplus (A_2 \times \{0\}) = \mathbb{F}_2^{13}$ *with* $|A_2 \times \{0\}| = 16$*, and* R_2 *has no spectrum. Hence T–S fails in an* elementary abelian 2*-group:* $\mathbb{F}_2^{13}$ *has order* $8192 = 2^{13}$ *and exponent* 2*, and carries a translational tile that is not spectral.*

Proof. The tiling is Lemma 6.1 with $A = A_2$ and $C_y = C_j$ for $y \in Y_j$, whose hypothesis is Proposition 7.2; so it costs the $17 \cdot 256$ element sums of that proposition, and $\mathbb{F}_2^{13}$ is never enumerated. For non-spectrality apply Lemma 6.2 with $D = D_2$ and $w = 12$: the hypothesis on D is the last clause of Proposition 7.3 and the hypothesis on w is Theorem 7.5. Every family of characters pairwise orthogonal on R_2 therefore has at most $|L_2| \cdot 12 = 32 \cdot 12 = 384$ elements, and $384 < 512 = |R_2|$. That the ambient group is elementary abelian of order 2^{13} is visible in the group itself, not a hypothesis of the statement. □

7.3. How far R_2 is from spectral.

Theorem 7.7. *Let* $\Theta \subset \mathbb{F}_2^8$ *be the* 12*-point set of Theorem 7.5 and put* $\Lambda_{384} = \Theta \times L_2 \subseteq \mathbb{F}_2^{13}$*. Then* $|\Lambda_{384}| = 384$ *and the characters* χ_u*,* $u \in \Lambda_{384}$*, are pairwise orthogonal on* R_2*:*

$$\sum_{a \in R_2} \chi_u(a)\,\overline{\chi_v(a)} = 0 \qquad \text{for all distinct } u, v \in \Lambda_{384}.$$

Moreover no *family of characters pairwise orthogonal on R_2 has more than* 384 *elements. A spectrum would need* $|R_2| = 512$.

Proof sketch. The upper bound is Lemma 6.2 with $D = D_2$, $w = 12$ and $|L_2| = 32$, exactly as in the proof of Theorem 7.6; it rests on the same computation as Theorem 7.5 and on nothing further. For the lower bound, the $12 \cdot 12 \cdot 32$ differences $(c - c', y)$ with $c, c' \in \Theta$ and $y \in L_2$ are checked to be zeros of $\widehat{1_{R_2}}$; *that check — 4,608 vanishing tests — is the exhaustive part.* The pairwise orthogonality is then obtained from it as an identity between complex sums, through $\chi_u \overline{\chi_v} = \chi_{u-v}$ and the exact vanishing test, and not merely as a statement about fibre counts.

[1] states only that R_2 is not spectral and computes no such quantity. The witness is not an accident of a search: it is the maximum clique of Theorem 7.5 crossed with the whole layer group, and the proof of Lemma 6.2 shows that a family attaining the bound must carry a maximum clique of $\mathrm{Cay}(X, D)$ in every one of its $|Y|$ fibres. ☐

8. A smaller counterexample: 128 points in $\mathbb{F}_2^{11}$

[1] closes with the following, which is the starting point of this section:

> "The constructions do not determine the least possible ambient orders or dimensions. In particular, the numbers 13 and 9 arise from the present base-plus-layer certificates rather than from lower-bound theorems.
>
> The fiber–clique lemma suggests two routes to smaller examples. One may reduce the dimension of the base group while retaining enough tiling complements, or reduce the number of layer types while forcing the weighted horizontal Fourier zero graph to have clique number below the complement size. Both are finite but structurally constrained search problems."

We carry out the second route on the authors' own base data. The observation that makes it a small search is visible in Lemma 6.2: the inequality that has to hold, $\omega(\mathrm{Cay}(X, D)) < q$, involves neither $|Y|$ nor $|R|$. Shrinking the layer group therefore costs nothing on the obstruction side, while it shrinks the tile and the ambient group. Moreover W, and hence D, depends only on the *multiset* of complements assigned to the layers, by the observation [1] records in its Remark 4.2, immediately after that lemma. So with the base $A_2 \subset \mathbb{F}_2^8$ and its seventeen printed complements fixed, and with $Y = \mathbb{F}_2^k$, the search space at layer dimension k is the set of multisets of size 2^k drawn from seventeen items: 17, 153, 4,845 and 735,471 of them for $k = 0, 1, 2, 3$. Exactly three of the 735,471 at $k = 3$ have $\omega < 16$, and none of the smaller ones does (Remark 8.4). The following is one of the three.

Theorem 8.1. *For $y \in \mathbb{F}_2^3$ with integer code $i \in \{0, \dots, 7\}$ let $j(y)$ be the entry in position i, counting from* 0*, of the list*

$$(\,2,\ 6,\ 9,\ 10,\ 12,\ 13,\ 14,\ 16\,),$$

let $C_{j(y)}$ be the corresponding complement of Appendix B, and put

$$R_S = \bigcup_{y \in \mathbb{F}_2^3} \big(C_{j(y)} \times \{y\}\big) \ \subseteq\ \mathbb{F}_2^8 \times \mathbb{F}_2^3 = \mathbb{F}_2^{11}.$$

Then:

(i) $\mathbb{F}_2^{11}$ *is elementary abelian of order* 2048*;*

(ii) $|R_S| = 128$ *and* $R_S \oplus (A_2 \times \{0\}) = \mathbb{F}_2^{11}$, *with* $|A_2 \times \{0\}| = 16$*;*

(iii) *the set D_S of nonzero $v \in \mathbb{F}_2^8$ for which $(v, 0)$ is a zero of $\widehat{1_{R_S}}$ has* 118 *elements, and the clique number of* $\mathrm{Cay}(\mathbb{F}_2^8, D_S)$ *is exactly* 12*, attained by the same* 12*-point set* Θ *as in Theorem 7.5;*

(iv) *R_S has no spectrum. Indeed $\Theta \times \mathbb{F}_2^3$ is a family of exactly* 96 *characters of $\mathbb{F}_2^{11}$ that are pairwise orthogonal on R_S, no family of characters pairwise orthogonal on R_S has more than* 96 *elements, and a spectrum would need* $|R_S| = 128$.

Proof sketch. (i) is the group. (ii) is Lemma 6.1 applied to the eight complements named, whose hypothesis is eight of the seventeen cases of Proposition 7.2; the cardinality is $8 \cdot 16 = 128$ by the same lemma.

(iii): $|D_S| = 118$ is decided from the list, and that the list is exactly the set of nonzero horizontal Fourier zeros of R_S is an exhaustive check over the 256 frequencies of $\mathbb{F}_2^8$. That Θ is a clique of $\mathrm{Cay}(\mathbb{F}_2^8, D_S)$ is a check of its $\binom{12}{2}$ differences. For the upper bound, the same translation argument as in Theorem 7.5 reduces the question to "does the 118-point set D_S contain a 12-clique?", and *that is the claim resting on exhaustive computation*: the same plain increasing-index enumeration, with the same completeness proof and again with no colouring bound, returns *false* in 52,495 nodes and 320,574 adjacency tests as measured outside the formal development. At target 11 it returns *true*, in 110 nodes (Proposition 10.2(v)).

(iv) is Lemma 6.2 with $X = \mathbb{F}_2^8$, $Y = \mathbb{F}_2^3$, $D = D_S$ and $w = 12$, giving $8 \cdot 12 = 96$; and $96 < 128$, so there is no spectrum. The witness $\Theta \times \mathbb{F}_2^3$ has $12 \cdot 8 = 96$ elements and its $12 \cdot 12 \cdot 8 = 1{,}152$ differences are checked to be zeros of $\widehat{1_{R_S}}$, from which the pairwise orthogonality follows as an identity between complex sums.

The character theory of $\mathbb{F}_2^{11}$ is the same ± 1 theory as that of $\mathbb{F}_2^{13}$ at a different layer group, and is redone for this group — including the nondegeneracy of its dot product, checked as in Proposition 7.1, so that a subset of $\mathbb{F}_2^{11}$ again names that many distinct characters; Lemmas 6.1 and 6.2, being stated for arbitrary finite abelian X and Y, are reused unchanged, which is what makes the second example cheap. □

Corollary 8.2. *There is a subset of the elementary abelian group* $\mathbb{F}_2^{11}$*, of order* $2048 = 2^{11}$*, that is a translational tile and is not spectral. Hence the conjecture of* [2, Conjecture 1.3]*, attributed there to* [3]*, already fails in an elementary abelian* 2*-group of order* 2048*. So the least* d *for which* $\mathbb{F}_2^d$ *carries a translational tile without a spectrum is at most* 11*, against the bound* 13 *that* [1] *gives; in its notation,* $d_2 \le 13$ *becomes* $d_2 \le 11$*.*

Proof. Theorem 8.1, items (i), (ii) and (iv). □

Remark 8.3 (the ratio 3/4; an observation). The three quantities computed above — 48 of 64 for Γ (Theorem 5.1 and Remark 5.2, the upper half measured), 384 of 512 for R_2 (Theorem 7.7) and 96 of 128 for R_S (Theorem 8.1) — happen to agree in one respect. The ratio is 3/4 in all three cases. For R_2 and R_S the common value 3/4 is $\omega/q = 12/16$ and is therefore forced by the two ingredients of Lemma 6.2 once both bounds are attained; for Γ, whose mechanism is the two-layer lemma and not the fibre–clique lemma, no such explanation is available and the coincidence is unexplained. We record it as an observation on three data points and claim nothing from it.

Remark 8.4 (what the layer search covered, and what it did not; computed outside the formal development). Everything in this remark is a computation reported here, not part of the formal development; only the single instance formalized in Theorem 8.1 is machine-checked. A branch-and-bound with a greedy-colouring bound, stopped at the first 16-clique, was run on every multiset of complements of the shapes described above, and the tables are complete:

k	layer group	ambient	$\lvert R\rvert$	multisets	with $\omega < 16$
0	trivial	$\mathbb{F}_2^8$	16	17	0
1	$\mathbb{F}_2^1$	$\mathbb{F}_2^9$	32	153	0
2	$\mathbb{F}_2^2$	$\mathbb{F}_2^{10}$	64	4,845	0
3	$\mathbb{F}_2^3$	$\mathbb{F}_2^{11}$	128	735,471	3

The $k = 3$ row took 4 minutes 49 seconds. Its three winners are the multisets of indices

$$\{4,6,9,10,12,13,14,16\}, \qquad \{2,6,9,10,12,13,14,16\}, \qquad \{3,7,9,10,12,13,14,16\},$$

with $|D| = 118$, 118 and 119 respectively; all three have $\omega = 12$, and the second is the one of Theorem 8.1. As a positive control the same program was run on the authors' own weight vector (3) and returned $|D| = 108$ and no 16-clique, reproducing Theorem 7.5.

These tables are complete only over the seventeen complements [1] prints. The full pool of tiling complements of A_2 in $\mathbb{F}_2^8$ was also enumerated — there are 193,536 of them — and every one of them is spectral, so no single tiling complement of A_2 is by itself a 16-point tile of $\mathbb{F}_2^8$ without a spectrum and the layered construction is doing real work. Over that full pool the case $k = 1$ — which would give a 32-point tile in $\mathbb{F}_2^9$, of order 512 — was scanned only under the filter "at most 4 cancellations of $\widehat{1_{C_a}} + \widehat{1_{C_b}}$ on the 147 frequencies where $\widehat{1_{A_2}}$ vanishes": none of the $\binom{193536}{2} \approx 1.87 \cdot 10^{10}$ pairs passes it. That is *not* a complete decision for $k = 1$: the pairs with five or more cancellations were never clique-tested, and testing them is about 2,000 core-days, which was not run. The case $k = 2$ over the full pool is $\binom{193536+3}{4} \approx 5.8 \cdot 10^{19}$ multisets and is out of reach by enumeration. Reducing the base dimension below 8 — the other route [1] names — would need a fresh exploratory search for a small tile in $\mathbb{F}_2^7$ with enough complements, and was not attempted. So 11 is an upper bound for the least dimension of an elementary abelian binary counterexample, and this paper does not claim it is the value.

9. The base tile is the bottleneck

The certificate of [1] presents D_2 as an intersection over its seventeen complements, and the search of Section 8 varies the complements. It is therefore worth knowing how much of D_2 the complements control. The answer is: none of it.

Lemma 9.1 (the tiling identity). *If $A \oplus C = \mathbb{F}_2^8$ then $\widehat{1_A}(v) \cdot \widehat{1_C}(v) = 0$ for every nonzero $v \in \mathbb{F}_2^8$.*

Proof. The map $(a, c) \mapsto a + c$ is a bijection $A \times C \to \mathbb{F}_2^8$, so

$$\widehat{1_A}(v)\,\widehat{1_C}(v) = \sum_{a\in A}\sum_{c\in C}(-1)^{\langle a+c,v\rangle} = \sum_{x\in\mathbb{F}_2^8}(-1)^{\langle x,v\rangle} = 0 \qquad (v \neq 0),$$

the last step being the vanishing of a nontrivial character sum, which is checked here by exhausting the 255 nonzero frequencies. Everything is exact integer arithmetic. □

Theorem 9.2. *$D_2 = \{v \neq 0 : \widehat{1_{A_2}}(v) \neq 0\}$; equivalently, $\widehat{1_{A_2}}$ vanishes at exactly 147 of the 255 nonzero frequencies of $\mathbb{F}_2^8$ and D_2 is the complementary 108. Consequently, let Y be any finite layer group and let C_y, $y \in Y$, be* any *tiling complements of A_2 in $\mathbb{F}_2^8$; if $D \subseteq \mathbb{F}_2^8$ contains every nonzero v at which $\sum_{y\in Y} \widehat{1_{C_y}}(v)$ vanishes, then $\mathrm{Cay}(\mathbb{F}_2^8, D)$ has a clique of 12 vertices, so $\omega(\mathrm{Cay}(\mathbb{F}_2^8, D)) \geq 12$.*

Proof. The first assertion is an exhaustive check over the 256 frequencies of $\mathbb{F}_2^8$, and the count 147 another; $255 - 147 = 108$ recovers the authors' $|D_2| = 108$. For the second, let $v \in D_2$. Then $v \neq 0$ and $\widehat{1_{A_2}}(v) \neq 0$, so Lemma 9.1 forces $\widehat{1_{C_y}}(v) = 0$ for every y, whence $\sum_y \widehat{1_{C_y}}(v) = 0$ and $v \in D$. Thus $D_2 \subseteq D$, and the 12-point set Θ of Theorem 7.5, all of whose differences lie in D_2, is a clique of $\mathrm{Cay}(\mathbb{F}_2^8, D)$. □

Corollary 9.3 (the authors' certificate is optimal for their base). *With base A_2, the fibre–clique method cannot produce a clique number below 12, whatever the layer group and whatever tiling complements are assigned to its layers; and the weights (3) of [1] achieve exactly 12. Against the complement size $q = 16$ the slack is therefore exactly 4, and it cannot be widened.*

Proof. Theorem 9.2 for the lower bound and Theorem 7.5 for the value. □

This is the sharp form of what the tables of Remark 8.4 found empirically: the failures at $k = 0, 1, 2$ are failures of the *extra* cancellations to stay few, never of the floor. It also isolates the contrast between the two examples of this paper. For the exponent-4 example the corresponding clique number is 16, equal to q exactly (Theorem 4.4), and the clique route is closed off altogether — which is why [1] needs its two-layer lemma there. For the binary example the clique route works, with slack 4; and by Corollary 9.3 that slack is all the room the base A_2 allows.

10. Controls

The statements above are of two shapes — "this set tiles" and "no set of this size does" — and both shapes can be made true by a misplaced quantifier. The following were checked to exclude that.

Proposition 10.1. (i) *The subgroup* $V = \mathbb{Z}_4^3 \times \{0\} \times \{0\} \times \{0\} \subset G$ *has* 64 *elements, tiles* G *with complement* $W = \{0\}^3 \times \mathbb{Z}_4 \times \mathbb{Z}_2 \times \mathbb{Z}_2$*, and is spectral, being its own spectrum. So neither property is vacuous, and both can hold at the size where* Γ *fails one of them.*

(ii) $K \times \{0\}$ *has exactly the cardinality a tiling complement of* Γ *needs,* $64 \cdot 16 = 1024$*, and is not one: the element* $\big((0001; 1), 0\big)$ *of* $G = H \times \mathbb{Z}_2$ *has no representation. So the tiling of Theorem 3.1 is not a consequence of counting.*

(iii) *No* 3*-point subset of* G *is a tile, since* $3 \nmid 1024$*.*

Item (ii) is the negative control that matters most here: the sufficient condition used in the proof of Theorem 3.1 needs the covering check as well as the cardinality identity, and (ii) exhibits a pair satisfying the identity and failing to tile. Lemma 3.3 doubles as a non-vacuity check for the Fourier test, with 240 witnesses of non-vanishing.

Proposition 10.2 (controls for the elementary abelian examples)**.** (i) *The subgroup* $V_2 = \mathbb{F}_2^8 \times \langle e_1 \rangle$ *of* $\mathbb{F}_2^{13}$ *has* 512 *elements, exactly* $|R_2|$*; it tiles* $\mathbb{F}_2^{13}$ *with complement* $\{0\} \times \langle e_2, \ldots, e_5 \rangle$*, and it is spectral, being its own spectrum. So neither property is vacuous here either, and both can hold at the size where* R_2 *fails one of them.*

(ii) *There is a* 16*-point set* $B \subset \mathbb{F}_2^{13}$ *with* $|R_2|\,|B| = 512 \cdot 16 = 8192 = |\mathbb{F}_2^{13}|$ *that is not a tiling complement of* R_2*: the element of code* 16 *has no representation. So the tiling of Theorem 7.6 is not a consequence of counting.*

(iii) *No* 3*-point subset of* $\mathbb{F}_2^{13}$ *or of* $\mathbb{F}_2^{11}$ *is a tile, since* 3 *divides neither* 8192 *nor* 2048*.*

(iv) *The Fourier test is not vacuous on either tile: some nonzero frequency is not a zero of* $\widehat{1_{R_2}}$*, and some nonzero frequency is not a zero of* $\widehat{1_{R_S}}$*.*

(v) *The clique enumeration behind Theorems 7.5 and 8.1 returns* true *at target* 11 *on both* D_2 *and* D_S*, in* 50 *and* 110 *nodes; so its* false *verdicts at target* 12 *are not the artefact of a procedure that always answers* false*.*

Item (v) is the control that matters most for the two clique theorems, since a clique enumeration that never succeeds would prove them both. Outside the formal development the layer search of Remark 8.4 carries its own positive control, described there: run on the authors' own weight vector it reproduces $|D_2| = 108$ and the absence of a 16-clique.

11. The printed certificates of [1], reproduced

The three items of Proposition 5.1 of [1] that Section 3 does not consume are reproduced here, because they are the certificate the authors print and because reproducing them is what confirms that the data used above is the data they published. Each is stated in the exact-integer form their verifier is described as using.

Proposition 11.1 ([1], Proposition 5.1(1))**.** E *is a spectrum of* K *in* H*:* $|E| = |K| = 16$ *and* $\widehat{1_K}(e' - e) = 0$ *for all distinct* $e, e' \in E$*; equivalently, the Gram matrix of the* 16×16 *matrix* $(i^{\langle k, e \rangle})_{k \in K,\, e \in E}$ *is* $16\, I_{16}$*.*

Proposition 11.2 ([1], Proposition 5.1(2))**.** $E \oplus T_0 = H$ *and* $E \oplus T_1 = H$*.*

Proposition 11.3 ([1], Proposition 5.1(4))**.** $|(K - K) \setminus \{0\}| = 152$*, and* $|\widehat{1_{T_0}}(d)|^2 \neq |\widehat{1_{T_1}}(d)|^2$ *for every one of those* 152 *elements* d*.*

The count 152 is the authors' printed `oriented_differences=152`; the equality of the diagonal Gram entries with $|K| = 16$ in Proposition 11.1 is immediate, and the off-diagonal entry at (e, e') is the Fourier coefficient of K at $e' - e$, so the integer condition stated there is the verifier's check rather than a weaker surrogate. In the formal development the integer form is transported to the complex-character statement along the layer-0 embedding $H \hookrightarrow G$, under which (1) and (2) agree.

Remark 11.4 (the node count, and how much of the branching rule it pins down). Of all the numbers [1] prints for this example, exactly one depends on an algorithm rather than on the data: `exact_cover_nodes=42` for the search of Proposition 5.1(3). An exact-cover search branching on the lowest uncovered element of H, with no normalisation, visits 426,203 nodes. Adding the authors' own normalisation — a hypothetical complement is translated to contain 0 — brings this to 24,478, and that is the tree replayed inside the formal development. Both figures are robust: two independent implementations, differing in the order in which the 16 candidate translates are tried, return them unchanged. The printed 42 is not robust in the same way. It is reached by adding, on top of the normalisation, a fewest-options-first branching rule — resolve the uncovered element with the fewest admissible translates — in one particular implementation: the one that scans the uncovered elements in the fixed order of H, keeps the first minimiser, and stops the scan as soon as it meets an element with at most one admissible translate. Scanning to a true minimiser instead gives 33; keeping the last minimiser rather than the first gives 50. So the printed count places the authors' search in the fewest-options-first family without determining it, which is still more than a paper that prints no count tells a re-checker. The remaining printed values for this example,

```
K_spectral=1  K_tile=0  E_plus_T0=H  E_plus_T1=H
Fourier_magnitude_separated=1  oriented_differences=152
Gamma_tile=1  Gamma_spectral=0  ambient=C4^4xC2^2  size=64,
```

depend only on the data, and are reproduced by Propositions 2.4, 11.1, 11.2, 11.3 and Theorems 3.1, 3.2, 3.4.

Remark 11.5 (the printed output for the binary example). For R_2 the authors print

```
VERIFIED_ELEMENTARY_TILE_NONSPECTRAL
group=F2^13  group_order=8192  tile_size=512
tiling_complement_size=16  common_fourier_zeros=108
common_zero_clique_number=12  no_16_clique=1
weighted_noncancellation=1  minimum_abs_weighted_fourier=4.
```

Every value in this block depends on the coordinate data alone, and each is reproduced above: the first two lines by Proposition 7.1 and Theorem 7.6, the count 108 by Proposition 7.3, the clique number and `no_16_clique=1` by Theorem 7.5 (which gives $12 < 16$), and the last line by Proposition 7.4. Unlike the exponent-4 block, this one prints no node count, so nothing in it depends on the authors' branching rule; the figure of 425 search nodes appears in the proof of their Proposition 6.1 rather than in the output, and it is a property of their colouring branch-and-bound, not of the data. The plain enumeration replayed in Theorem 7.5, which carries a completeness proof and no colouring bound, needs 26,683 nodes for the same conclusion.

Remark 11.6 (availability of the verifier programs). [1] names three verifier programs and prints their output. As of 3 September 2026 the arXiv record of [1] lists no ancillary files, so those programs do not appear to be attached to it; what makes the paper reproducible is the coordinate data in its appendices, which is what is consumed here.

Remark 11.7 (two bibliographic slips in [1]). Neither affects any mathematical claim, and both are recorded only so that a reader following the references is not delayed. First, the reference to Kolountzakis and Matolcsi [10] is given as *Collect. Math.* ***57*** *(2006), 281–291*, with no indication that the paper is in the *Extra* volume of that year; EUDML records it as "Volume: 57, Issue: Extra, page 281-291", and *zbMATH* files it outside the three regular issues of volume 57. The Extra volume is paginated independently of the regular one, whose third issue of 2006 runs past page 340; so page 281 of volume 57 exists and belongs to a different article, and the citation as printed resolves to the wrong place rather than to nothing. Second, the reference to Fallon, Mayeli and Villano [9] is given as *Proc. Amer. Math. Soc.* ***147*** *(2019), 4197–4207*; when checked on 3 September 2026 the DOI `10.1090/proc/14750` still resolved to the AMS *Early View* listing, where the article carries no issue and neither a start nor an end page, and Crossref likewise records no volume and no page range. So the volume and pages appear to be anticipated rather than assigned.

Remark 11.8 (related work not cited in [1]). Zhang [11] exhibits, six weeks before [1], two 60-point subsets of $\mathbb{Z}_{60} \times \mathbb{Z}_{12}$ showing that both directions of Fuglede's conjecture fail in that group. Since $|\mathbb{Z}_{60} \times \mathbb{Z}_{12}| = 720 = 2^4 \cdot 3^2 \cdot 5$ is not a prime power, that result leaves the p-group conjecture untouched and there is no overlap of content; we mention it because the two papers share a theme and a reader may expect to see them together.

Changes from version 1 of this paper. Nothing stated in version 1 is withdrawn. Every lemma, proposition, theorem and corollary of that version stands here with its statement unchanged, and the two claims that version narrowed during refereeing — the derivation of $|S| \leq 32$ from the clique–coclique bound applied to the 16-clique K rather than from the conditional of Proposition 4.2, and the account in Remark 11.4 of how much of the authors' branching rule their printed node count pins down — are kept verbatim. What version 1 called "the counterexample" is now called the exponent-4 example, since a second one is treated.

Everything else is new: Sections 6–9 in their entirety, together with Proposition 10.2, Remark 11.5, Appendix B, items (a) and (b) of Section 1, the corresponding paragraphs of the abstract, and the second half of Section 12. Three parts of version 1 acquire company rather than corrections: item (c) of Section 1, which reported the orthogonality defect of Γ alone and now reports three; item (e), which reported the reproduction of one printed certificate and now reports two; and Section 10, whose controls are now stated for all three tiles. Item (d) is version 1's own account of the exponent-4 search, reworded only where it had to name which example it is about. The title and the abstract were rewritten to match; the paragraph of Section 1 on what is not claimed gained two sentences, one on what is new about the smaller example and one on the least dimension of an elementary abelian counterexample; and the paragraph before it gained the reading of [1]'s undefined symbol d_p that the last of those sentences uses.

12. Verification

Every numbered lemma, proposition, theorem and corollary of Sections 2–11 has been formally verified in Lean 4 (toolchain `v4.32.0`) against *Mathlib* [13, 14], with three qualifications: Corollaries 4.5 and 9.3 are not separate formal statements but arithmetic rephrasings of formal ones — the first of the clique–coclique bound (Lemma 4.1) applied to the formal 16-clique of Proposition 4.2, together with Theorem 4.4; the second of Theorems 9.2 and 7.5. Corollary 8.2 is formal in its first two sentences, exactly as Corollary 3.5 is; its last sentence, which records the consequence for the least dimension and restates that consequence in the notation of [1], is not. The remarks are outside the formal development by design: Remarks 4.6 and 5.2 report the two measured clique numbers of the exponent-4 example, Remark 8.4 reports the layer tables and the pool scans, Remark 11.4 reports node counts, Remark 8.3 is an observation about three numbers each of which is proved elsewhere, and Remarks 11.5–11.8 are accounting or bibliography.

The development is thirteen modules and 3,011 lines — counting every line of every file, the convention under which version 1 reported 1,424 for its six — and it contains no `sorry` and declares no axiom by hand. Six modules carry the exponent-4 example and rebuild cleanly in 247 s of wall time at a peak of about 7.9 GB resident, of which the *Mathlib* import is the bulk. Seven further modules, in a namespace of their own, carry the elementary abelian side; a clean sequential rebuild in import order takes 592 s of wall time — 16, 26, 122, 49, 164, 188 and 27 seconds — at a peak of 6.6 to 8.0 GB, again dominated by the import. The general fibre–clique lemmas (Lemmas 6.1 and 6.2) are stated once, for arbitrary finite abelian X and Y, and are used unchanged by both elementary abelian examples; the same is true of the step that translates a clique of a Cayley graph so as to contain 0, which is what makes the second example cheap.

Every reasoning step is free of compiled evaluation and depends on no axioms beyond `propext`, `Classical.choice` and `Quot.sound`: this includes the exact vanishing tests, both criteria (Lemmas 2.2 and 2.3), the clique–coclique bound (Lemma 4.1), the layer argument of Theorem 3.4, the translation normalisation and the completeness of the exact-cover search of Theorem 3.2, the completeness of the clique enumeration used in Theorems 4.3, 7.5 and 8.1 and the translation step in front of it, the fibre tiling lemma and the quantified fibre–clique bound, the tiling identity of Lemma 9.1, the deduction of Theorem 9.2 from it, and the transport of every integer Fourier test to the corresponding statement about complex characters. What is delegated to compiled evaluation, through Lean's `native_decide`, is the finite data checks and the replayed searches, listed next; a few of the smallest data checks — among them the two cardinalities $|D_2| = 108$ and $|D_S| = 118$ and the two verifications that Θ is a clique — are reduced by the Lean kernel itself instead, and are then free of compiled evaluation as well. For the exponent-4 example these are the nondegeneracy of (2) (Proposition 2.4); the covering checks behind Theorem 3.1 and Proposition 11.2; the 240 non-vanishing tests of Lemma 3.3 and the 240 membership tests of Proposition 4.2; the 1128 vanishing tests of Theorem 5.1; the Gram check of Proposition 11.1, the difference count and magnitude separation of Proposition 11.3, the controls of Proposition 10.1; and the two Boolean verdicts "the exact-cover search returns false" and "the clique enumeration returns false". For the elementary abelian examples they are the nondegeneracy of the dot products; the $17 \cdot 256$ covering checks of Proposition 7.2; the exhaustive checks behind Propositions 7.3 and 7.4; the identification of the listed 118-point set with the horizontal Fourier zero set of R_S; the 4,608 and 1,152 difference checks behind Theorems 7.7 and 8.1; the character-sum vanishing $\sum_x (-1)^{\langle x,v\rangle} = 0$, the identity $D_2 = \{v \neq 0 : \widehat{1_{A_2}}(v) \neq 0\}$ and the count 147 of Theorem 9.2; the controls of Proposition 10.2; and the two Boolean verdicts "no 12-clique inside D_2" and "no 12-clique inside D_S" together with their two positive controls at target 11. Each theorem depends on precisely the compiled checks its own proof uses — Theorem 3.2 on the exact-cover verdict alone, Theorems 4.3 and 4.4 on the clique verdict and the two transcription checks that identify the enumerated list with $(K - K) \setminus \{0\}$, Theorem 7.6 on the covering checks, the horizontal zero set of R_2 and the D_2 clique verdict, Theorem 8.1 additionally on the D_S clique verdict and the 1,152 difference checks, Theorem 9.2 on the character-sum vanishing and the two checks about $\widehat{1_{A_2}}$. No node count is an input to any proof, and none is verified: what each of the four searches — one exact cover and three clique enumerations — contributes is a Boolean, and their completeness is a theorem.

Three implementation notes, since they are the difference between a feasible check and an infeasible one. The state of the exact-cover and clique searches of the exponent-4 example is a 512-bit integer, and computing translate masks with 2^n rather than with a single left shift $1 \ll n$ made the exact-cover module fail to finish in 28 minutes, against 90 s after the change; and the normalisation of Theorem 3.2, which is the authors' own, is what brings the replayed tree from 426,203 nodes to 24,478. For the two binary clique theorems the corresponding remark is that *no colouring bound was needed at all*: the translation to 0 turns each of them into a search inside a set of 108 or 118 vertices, and the plain enumeration — the one whose completeness is proved — settles them in 26,683 and 52,495 nodes. That is why nothing about those searches has to be trusted, and it is also the reason the third example of [1], in $\mathbb{F}_3^9$, is not treated here: its clique step is "no 81-clique" on a 729-vertex graph, which the same colouring-free enumeration does not reach. Separately, every number quoted in this paper was first produced by an independent implementation, in exact integer arithmetic, that reads the coordinate data directly out of the LaTeX source of [1]; the formal development was written against that implementation afterwards, and the clique computations reported as measurements (Remarks 4.6, 5.2 and 8.4, and the cross-check in Theorem 4.3) were run with a separate branch-and-bound in C.

The source of the development is currently held in a private repository: repository reference to be added at submission.

Appendix A. The data

The sets $K, E, T_0, T_1 \subseteq H = \mathbb{Z}_4^4 \times \mathbb{Z}_2$ are those of Appendix A of [1], in its notation $(abcd; e)$; they are reproduced here because the results of Sections 4 and 5 are about them. The listing below was generated mechanically from the same machine-readable transcription that the formal development uses.

$$K = \{(0000;0),\ (0200;0),\ (0000;1),\ (1000;0),\ (3202;0),\ (0100;0),\ (0010;0),\ (2112;1),\\ (2122;0),\ (0220;1),\ (0001;0),\ (3001;0),\ (1003;0),\ (2203;0),\ (0030;0),\ (2132;1)\},$$

$$E = \{(0000;0),\ (0012;0),\ (0221;0),\ (2231;0),\ (0203;1),\ (2213;1),\ (2020;1),\ (2032;1),\\ (1133;0),\ (2303;0),\ (2311;0),\ (3123;0),\ (0323;1),\ (0331;1),\ (1111;1),\ (3101;1)\}.$$

$$T_0 = \{(0000;0),\ (2002;0),\ (2301;1),\ (0303;1),\ (1301;0),\ (3303;0),\ (1320;0),\ (3322;0),\\ (0300;0),\ (2302;0),\ (0021;0),\ (2023;0),\ (1023;1),\ (3021;1),\ (0023;0),\ (2021;0),\\ (3323;1),\ (1321;1),\ (2321;0),\ (0323;0),\ (0002;0),\ (2000;0),\ (0302;0),\ (2300;0),\\ (1003;0),\ (3001;0),\ (1022;0),\ (3020;0),\ (3302;1),\ (1300;1),\ (3022;0),\ (1020;0)\},$$

$$T_1 = \{(0000;0),\ (3100;0),\ (0301;0),\ (1223;1),\ (1322;1),\ (0210;0),\ (0230;0),\ (1013;1),\\ (1033;1),\ (0200;1),\ (2123;1),\ (1023;0),\ (3121;1),\ (1123;1),\ (1301;0),\ (3303;0),\\ (2320;1),\ (0322;1),\ (0100;0),\ (2102;0),\ (1213;1),\ (1223;0),\ (1233;1),\ (3003;0),\\ (1203;0),\ (1001;0),\ (0222;1),\ (2220;1),\ (0000;1),\ (2002;1),\ (0002;0),\ (2000;0)\}.$$

The 24-point set $\Omega \subset H$ of Theorem 5.1, for which $\Omega \times \mathbb{Z}_2$ is a family of 48 characters pairwise orthogonal on Γ, is

$$\Omega = \{(0000;0),\ (0102;0),\ (0311;0),\ (0313;0),\ (1020;0),\ (1022;0),\ (1121;1),\ (1123;1),\\ (1231;0),\ (1233;0),\ (1330;1),\ (1332;1),\ (2101;1),\ (2103;1),\ (2212;1),\ (2310;1),\\ (3021;1),\ (3023;1),\ (3120;0),\ (3122;0),\ (3230;1),\ (3232;1),\ (3331;0),\ (3333;0)\}.$$

Appendix B. The binary data

The sets below are those of Appendix B of [1], in its notation: an element of $\mathbb{F}_2^8$ is written as the integer $\sum_{i<8} x_i 2^i$. They are reproduced here because Sections 7–9 are about them. The listing was generated mechanically from the same machine-readable transcription that the formal development uses.

The base tile is

$$A_2 = \{3, 22, 35, 53, 69, 74, 96, 101, 136, 144, 168, 179, 199, 204, 230, 231\} \subset \mathbb{F}_2^8,$$

and its seventeen tiling complements are:

j	C_j
0	0, 9, 36, 45, 48, 57, 72, 97, 150, 159, 199, 215, 222, 238, 243, 250
1	0, 9, 36, 45, 48, 57, 72, 81, 88, 97, 117, 124, 150, 159, 199, 238
2	0, 9, 36, 45, 48, 51, 57, 58, 72, 81, 88, 97, 199, 208, 217, 238
3	0, 9, 36, 45, 48, 51, 57, 58, 72, 97, 114, 123, 199, 238, 243, 250
4	0, 9, 36, 45, 48, 57, 72, 97, 181, 188, 199, 208, 215, 217, 222, 238
5	0, 9, 19, 36, 45, 58, 72, 97, 156, 181, 199, 208, 217, 238, 244, 253
6	0, 9, 19, 26, 36, 45, 114, 123, 133, 140, 178, 181, 187, 188, 199, 206
7	0, 9, 36, 39, 45, 46, 117, 124, 149, 150, 156, 159, 178, 187, 199, 206
8	0, 26, 40, 50, 82, 96, 104, 122, 142, 148, 166, 188, 198, 206, 220, 244
9	0, 40, 49, 57, 96, 104, 113, 121, 142, 151, 159, 166, 198, 206, 215, 223
10	0, 26, 39, 61, 71, 90, 96, 125, 129, 155, 166, 188, 198, 219, 225, 252
11	0, 25, 30, 39, 71, 94, 96, 121, 129, 166, 184, 191, 198, 223, 225, 248
12	0, 30, 44, 50, 96, 108, 114, 126, 138, 148, 166, 184, 198, 202, 212, 216
13	0, 8, 23, 31, 77, 82, 90, 101, 166, 174, 177, 185, 195, 235, 244, 252
14	0, 2, 28, 30, 71, 89, 101, 123, 164, 166, 184, 186, 195, 221, 225, 255
15	0, 20, 39, 51, 94, 106, 109, 121, 129, 149, 166, 178, 203, 204, 223, 248
16	0, 6, 18, 20, 88, 94, 106, 108, 160, 166, 178, 180, 202, 204, 248, 254

R_2 uses all seventeen, with the multiplicities (3); the smaller example R_S of Theorem 8.1 uses the eight with $j \in \{2, 6, 9, 10, 12, 13, 14, 16\}$, one on each layer of $\mathbb{F}_2^3$, in that order of the layer codes $0, \ldots, 7$.

The 12-point set $\Theta \subset \mathbb{F}_2^8$ of Theorem 7.5, a maximum clique of $\mathrm{Cay}(\mathbb{F}_2^8, D_2)$ and also of $\mathrm{Cay}(\mathbb{F}_2^8, D_S)$, is

$$\Theta = \{0, 3, 9, 16, 25, 68, 77, 93, 151, 152, 159, 199\}.$$

The two connection sets are not listed: D_2 is determined by A_2 alone (Theorem 9.2), and D_S is determined by the eight complements named above.

References

[1] S. Fan and M. Kadir, *Translational tiles without spectra in finite abelian p-groups*, arXiv:2609.00087v1 [math.CA], 31 August 2026. Read here from the author-supplied LaTeX source. Checked live on 3 September 2026: one version, primary category math.CA with no cross-list, no journal reference and no DOI beyond the arXiv one. Every number by which a statement or appendix of it is cited above is the one v1 carries, resolved by compiling that source rather than read off the text: Theorem 1.2; Lemma 3.1, Proposition 5.1 and Appendix A for the exponent-4 example; Lemma 4.1, Remark 4.2, Proposition 6.1 and Appendix B for the binary one.

[2] R. D. Malikiosis, *A linear programming approach to Fuglede's conjecture in $\mathbb{Z}_p^3$*, Sampling Theory, Signal Processing, and Data Analysis **22** (2024), article 3, doi:10.1007/s43670-023-00078-2; arXiv:2212.13592. Conjecture 1.3 is the statement refuted by [1]; the quotation and the number 1.3 were both read off the LaTeX source of arXiv:2212.13592v1, where the conjecture is the third statement of §1 on the counter shared by its theorem-like environments.

[3] R. Shi, *Fuglede's conjecture holds on cyclic groups $\mathbb{Z}_{pqr}$*, Discrete Anal. (2019), Paper No. 14, 14 pp., doi:10.19086/da.10570; arXiv:1805.11261. Title, year and paper number confirmed against the journal's own record (published 15 October 2019, cited there as *Discrete Analysis* 2019:14).

[4] B. Fuglede, *Commuting self-adjoint partial differential operators and a group theoretic problem*, J. Funct. Anal. **16** (1974), no. 1, 101–121, doi:10.1016/0022-1236(74)90072-X.

[5] T. Tao, *Fuglede's conjecture is false in 5 and higher dimensions*, Math. Res. Lett. **11** (2004), no. 2–3, 251–258, doi:10.4310/mrl.2004.v11.n2.a8.

[6] B. Farkas, M. Matolcsi and P. Móra, *On Fuglede's conjecture and the existence of universal spectra*, J. Fourier Anal. Appl. **12** (2006), no. 5, 483–494, doi:10.1007/s00041-005-5069-7. The tile without a spectrum in $\mathbb{Z}_{24}^3$; that group is not a p-group.

[7] A. Iosevich, A. Mayeli and J. Pakianathan, *The Fuglede conjecture holds in $\mathbb{Z}_p \times \mathbb{Z}_p$*, Anal. PDE **10** (2017), no. 4, 757–764, doi:10.2140/apde.2017.10.757.

[8] S. J. Ferguson and N. Sothanaphan, *Fuglede's conjecture fails in 4 dimensions over odd prime fields*, Discrete Math. **343** (2020), no. 1, 111507, doi:10.1016/j.disc.2019.04.026; arXiv:1901.08734.

[9] T. Fallon, A. Mayeli and D. Villano, *The Fuglede conjecture holds in $\mathbb{F}_p^3$ for $p = 5, 7$*, Proc. Amer. Math. Soc., *Early View*; doi:10.1090/proc/14750; arXiv:1902.02936. See Remark 11.7: when checked on 3 September 2026 the article still carried no volume and no page range.

[10] M. N. Kolountzakis and M. Matolcsi, *Complex Hadamard matrices and the spectral set conjecture*, Collect. Math. **57** (2006), Extra volume, 281–291. See Remark 11.7.

[11] T. Zhang, *Both directions of Fuglede's conjecture fail in dimension two*, arXiv:2607.15632v1, 17 July 2026. Checked live on 3 September 2026: one version, not withdrawn, no journal reference.

[12] C. Godsil and G. Royle, *Algebraic Graph Theory*, Graduate Texts in Mathematics **207**, Springer, New York, 2001, doi:10.1007/978-1-4613-0163-9. Cited as a general reference for the vertex-transitive clique–coclique inequality; no chapter or theorem number of it is asserted here, and the form used in the text (Lemma 4.1) is proved in place.

[13] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction — CADE 28, Lecture Notes in Computer Science **12699**, Springer, 2021, 625–635, doi:10.1007/978-3-030-79876-5_37.

[14] The mathlib Community, *The Lean mathematical library*, in: Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020), ACM, 2020, 367–381, doi:10.1145/3372885.3373824.