NON-UNIQUE-PRODUCT SETS IN THE PROMISLOW GROUP: A CENTRING THEOREM, AND BOUNDS BY WORD DIAMETER

KOREA SUPERINTELLIGENCE LABS

Abstract. A finite subset A of a group is *non-UP* if no element of $A \cdot A$ is a product ab with $a, b \in A$ in exactly one way. In the Promislow group P — the orientable Hantzsche–Wendt Bieberbach group of dimension three, where Promislow found a 14-element non-UP set in 1988 and where Gardam refuted the unit conjecture in 2021 — the least size of a non-UP set is not known; it lies between 8, the bound valid in every torsion-free group, and 14. What has narrowed the question is a statement about a ball: the minimum over non-UP subsets of the word-metric ball $B(r)$ is 14 for $3 \leq r \leq 6$, and a ball computation cannot be promoted to a statement about the group, because the one-set non-UP condition is not translation invariant.

We give a normal form that is a conjugation rather than a translation, and therefore is a symmetry of the problem. For each of the three coordinate directions r let σ_r be the diagonal character of the point group, $N_r = \ker \sigma_r$ the corresponding index-two subgroup, and Ψ_r the r-th translation coordinate. We prove that a non-UP set A satisfies $\max \Psi_r(A \cap N_r) = -\min \Psi_r(A \cap N_r) =: h_r \geq 0$ — an *absolute* constraint, with no translation freedom left — that $\Psi_r(A \setminus N_r)$ occupies a window of width exactly $2h_r$, that A meets at least three of the four point-group fibres, and consequently that some conjugate of A by a translation lies in an explicit box $U(h)$ determined by $h = (h_1, h_2, h_3)$ alone. A search of $U(h)$ is then a statement about all of P.

Sweeping h over finite ranges gives three such statements: no non-UP subset of P, of any cardinality, has word diameter at most 3; no non-UP subset of P of cardinality at most 13 has word diameter at most 5; and a non-UP subset of P whose three spreads are all at most 2 has spread vector a cyclic rotation of Promislow's own $(1, 2, 1)$ and at least 14 elements. The last is a complete classification of 27 profiles. All three are sharp against Promislow's witness, which has spread vector $(1, 2, 1)$ and word diameter exactly 6. Every theorem below is machine-checked; the underlying refutations are 43 unsatisfiability results, 13 discharged inside the proof kernel and 30 by replaying clausal proofs.

1. Introduction

A nonempty finite subset A of a group G has the *unique product* property if some $w \in A \cdot A = \{ab : a, b \in A\}$ equals ab for exactly one ordered pair $(a, b) \in A \times A$; if no such w exists we call A *non-UP*. More generally a pair (A, B) of nonempty finite sets is non-UP if every element of $A \cdot B$ has at least two representations ab with $a \in A$, $b \in B$. Groups in which no finite subset is non-UP are the *unique product groups*; if G is one, then for every field $\mathbb{K}$ the group ring $\mathbb{K}[G]$ has no zero divisors and only trivial units. So non-UP sets are the combinatorial obstruction standing between Kaplansky's conjectures and the torsion-free hypothesis.

The standard torsion-free example of a group that is not a unique product group is Promislow's [7]: a 14-element non-UP set inside the orientable Hantzsche–Wendt

Bieberbach group P of dimension three, the fundamental group of the Hantzsche–Wendt flat 3-manifold. The same group carries Gardam's counterexample to the unit conjecture [4], which is why its non-UP sets are of interest beyond the example itself. Write

$$m_1(P) = \min\{|A| : A \subseteq P \text{ finite non-UP}\},$$
$$m_2(P) = \min\{|A| + |B| : (A, B) \text{ a non-UP pair in } P\}.$$

Nielsen and Soelberg [6] prove $m_1 \geq 8$ and $m_2 \geq 16$ in every torsion-free group, and Promislow's set gives $m_1(P) \leq 14$.

The source, and the shape of what it leaves open. Tabei [9] computes the *ball-limited* minima. Measuring length by the word metric of $\{x^{\pm 1}, y^{\pm 1}\}$ for the two standard generators and writing $B(r)$ for the ball of radius r about the identity, [9, Theorem 3.1] states that $r = 3$ is the least radius for which $B(r)$ contains a non-UP set at all, and that for $3 \leq r \leq 6$ the minimum cardinality of a non-UP subset of $B(r)$ is exactly 14. In the two-sided problem $m_2(P) \leq 24$ by an explicit pair inside $B(3)$, and 24 is the minimum total over pairs inside $B(3)$ [9, Proposition 3.3]. Both global questions — is $m_1(P) = 14$? is $m_2(P) = 24$? — remain open; the first is [9, Question 3.2].

The obstruction to converting the ball results into global ones is stated in the source and is not a matter of computing power. Two-sided non-UP-ness is invariant under $(A, B) \mapsto (gA, Bh)$, so a two-sided search may normalise; the one-set condition $A \cdot A$ has no such invariance, and indeed every nontrivial translate of Promislow's 14-set inside its own ball acquires a unique product. A hypothetical small non-UP set of small diameter sitting far from the identity therefore cannot be carried into $B(6)$, and [9, §5], titled *Why ordering does not bound the diameter*, argues that no ordering or convexity argument supplies the missing bound: "A diameter bound for minimal non-UP sets in P, if one exists, must use the affine geometry more globally."

The source does give one effective route around the obstruction: [9, Theorem 6.2] re-realises a hypothetical non-UP n-set — rather than translating it — inside the ball of radius $D(n) \leq 4^n \operatorname{poly}(n)$, so Question 3.2 is decidable in principle, and [9, Conjecture 6.4] conjectures that the true bound is $O(n^{1/3})$. The bound as proved is far outside search range, and it is a bound on a radius, so it is again a ball that would have to be searched.

What is settled here. This paper replaces the ball by a box determined by the set's own invariants, and replaces translation by conjugation, which *is* a symmetry of the one-set problem.

Section 3 is the mathematics, and contains no computation. For $r \in \{1, 2, 3\}$ let $\sigma_r \colon P \to \{\pm 1\}$ read off the r-th diagonal entry of the point part, let $N_r = \ker \sigma_r$ — an index-two subgroup — and let Ψ_r be the r-th translation coordinate, doubled so as to be an integer. These satisfy a cocycle identity $\Psi_r(gh) = \Psi_r(g) + \sigma_r(g)\Psi_r(h)$, from which everything follows. Theorem 5 shows that a non-UP set A meets both cosets of N_r and satisfies

(E1) $\max \Psi_r(A \cap N_r) = -\min \Psi_r(A \cap N_r) =: h_r \ \geq 0$,

(E2) $\Psi_r(A \setminus N_r)$ occupies a window of width $2h_r$.

Equation (E1) is an absolute constraint: it says the part of A on which $\sigma_r = +1$ is exactly centred at the origin in coordinate r, and it is precisely what fails for translates. Corollary 6 records that A meets at least three of the four point-group fibres, where [9, Corollary 4.6] proves that it meets at least two. Theorem 7, the *centring theorem*, uses (E2) and the fact that conjugation by a translation shifts Ψ_r off N_r by an arbitrary even integer, and fixes it on N_r, to place a conjugate of A inside the explicit box

$$U(h) = \big\{g \in P:\ |\Psi_r(g)| \le h_r \text{ for every } r \text{ with } g \in N_r,\\ \text{and } |\Psi_r(g)| \le h_r + 1 \text{ for every } r \text{ with } g \notin N_r\big\}.$$

Nothing about the location of A enters. Finally, if A has word diameter at most d then $2h_r \le \max\{\Psi_r(g) : \ell(g) \le d\}$, which converts a diameter hypothesis into a finite range of h (Proposition 8).

Sections 4 and 5 turn this into decisions. Refuting "there is a non-UP set of size at most k inside $U(h)$ with spread vector exactly h" for *every* h in a range refutes the existence of such a set anywhere in P with spread vector in that range. Three sweeps are carried out:

- $h \le (1,1,1)$, no cardinality bound, 8 instances: every non-UP subset of P has $h_r \ge 2$ for some r, and none has word diameter at most 3 (Theorem 10);
- $h \le (2,2,1)$ with $|A| \le 13$, 18 instances: no non-UP subset of P of cardinality at most 13 has word diameter at most 5; equivalently every non-UP set of word diameter at most 5 has at least 14 elements (Theorem 11);
- $h \le (2,2,2)$, no cardinality bound, 27 instances: exactly three spread profiles with all coordinates at most 2 carry a non-UP set, namely the cyclic orbit $(1,1,2), (1,2,1), (2,1,1)$ of Promislow's own profile, and at those three the size is at least 14 (Theorem 12).

The third is a complete classification of a 27-cell table, not a bound. All three speak about sets sitting anywhere in P, which is what the ball statements of [9] cannot do; the frontier variable becomes the word *diameter*, where in the ball picture it was the radius. For scale, the box $U(2,2,2)$ has 99 elements while $|B(6)| = 239$ and $|B(7)| = 363$; the three sweeps together retain about 8.8 megabytes of certificate, the largest single file being 3.8 megabytes, against 2.3 megabytes for the one ball statement at $r = 4$ that this development also carries.

The sharpness is Promislow's own set: its spread vector is $(1,2,1)$ and its word diameter is exactly 6 (Proposition 14), so the diameter bound 5 of Theorem 11 is one step below the known witness, the bound $\max_r h_r \ge 2$ of Theorem 10 is attained, and the classification of Theorem 12 contains the profile that is realised.

What is not settled. Nothing here bounds the diameter or the spread of a hypothetical minimal non-UP set. The centring theorem converts a diameter *hypothesis* into a finite box; it does not produce a diameter bound, and [9, Question 3.2] is untouched. What changes is the shape of the remaining question: it becomes "is $\max_r h_r$ bounded over non-UP sets of size $8 \le n \le 13$ in P?", and each fixed value of that bound is now decidable by a search over a box rather than over a ball. The two-sided minimum $m_2(P) \in [16, 24]$ is also untouched; see Remark 17.

Provenance. The peak lemma of §3 (Lemma 4) is not new in substance. Bowditch [1, §6], in the section where he shows that this group is not diffuse, observes that it has three index-two subgroups, each an extension of $\mathbb{Z}^2$ by $\mathbb{Z}$ and therefore diffuse. Those three subgroups are N_1, N_2, N_3. Indeed his i-th one is generated by $\{T_i, T_{i+1}^2, T_{i+2}^2\}$, where T_i is his glide translation with axis parallel to the i-th coordinate direction, so that $\sigma_i(T_i) = +1$; in our coordinates $T_1 = x$, $T_2 = y$ and $T_3 = (xy)^{-1}$. The squares are translations, so $\sigma_i = +1$ on all three generators and that subgroup is contained in N_i; both have index two, so they coincide. Combined with his Proposition 1.1, that a diffuse group has the two-unique-products property, and with a one-line coset reduction, this gives the existence half of Lemma 4. What we use, and did not find in that literature, is the *located* form: the uniquely represented product is a named pair realising the maximum of Ψ_r, and that is what makes the comparison of blocks in the proof of Theorem 5 possible at all. The source proves the single-fibre case ([9, Lemma 4.5]) and calls it folklore. The remaining statements of §3 — (E1), (E2), the three-fibre corollary and the centring theorem — and the three sweeps of §5 we believe to be new. The searches behind that belief are described in §7. We did not read the paywalled [6], Soelberg's thesis, or the literature on diffuseness of Bieberbach groups in full, so the negative should be read as "not found", not as "new".

2. The group, and two coordinate functions

The model. Following [9, §2] we realise P as a group of affine isometries $w \mapsto Mw+v$ of $\mathbb{R}^3$, where M runs over the Klein four-group $K = \{I, X, Y, Z\}$ of diagonal sign matrices of determinant $+1$,

$$X = \mathrm{diag}(1, -1, -1), \qquad Y = \mathrm{diag}(-1, 1, -1), \qquad Z = \mathrm{diag}(-1, -1, 1),$$

and $v \in \frac{1}{2}\mathbb{Z}^3$. We store elements as pairs (M, t) with $t = 2v \in \mathbb{Z}^3$, so that all arithmetic is exact; the group law and inversion read

$$(M, t)(M', t') = (MM',\, t + Mt'), \qquad (M, t)^{-1} = (M, -Mt).$$

Membership in P is the parity condition $t \equiv \varepsilon(M) \pmod 2$, where $\varepsilon(I) = (0, 0, 0)$, $\varepsilon(X) = (1, 1, 0)$, $\varepsilon(Y) = (0, 1, 1)$ and $\varepsilon(Z) = (1, 0, 1)$. The generators are $x = (X, (1, 1, 0))$ and $y = (Y, (0, 1, 1))$, so that x^2, y^2 and $(xy)^2$ are the translations $(2, 0, 0)$, $(0, 2, 0)$ and $(0, 0, -2)$ in doubled coordinates, and the two Hantzsche–Wendt relations $x^{-1}y^2x = y^{-2}$ and $y^{-1}x^2y = x^{-2}$ of Promislow's presentation hold. Write $\pi\colon P \to K$ for the point-group projection.

Proposition 1 (the model)**.** *The parity condition defines exactly the subgroup $\langle x, y\rangle$ of the ambient group, both inclusions. This group is torsion-free. For all u, v and all finite A, B, the pair (uA, Bv) is non-UP if and only if (A, B) is.*

The last clause is the two-sided translation invariance; it has no one-set analogue, and that is the whole difficulty. Torsion-freeness is [9, Remark 2.1]: an element is either a translation, of infinite order, or has point part of order 2, in which case its square is a nonzero translation because the relevant coordinate of v is a half-integer.

Word metric and the published data. $B(r)$ is the ball of radius r about the identity in the word metric of $\{x^{\pm1}, y^{\pm1}\}$, and $\ell(g)$ is word length. We say A has *word diameter at most d*, written $\mathrm{diam}(A) \le d$, if $a^{-1}a' \in B(d)$ for all $a, a' \in A$. Unlike containment in a ball, this condition is invariant under both translation and conjugation.

Proposition 2 (published data, reverified)**.** *The ball sizes are*

$$|B(1)|, \dots, |B(7)| \;=\; 5,\ 17,\ 41,\ 83,\ 147,\ 239,\ 363.$$

The set A_0 of [9, §3.1]*, given by the reduced words*

$$\begin{aligned}&\{x,\ x^{-1},\ xy^2,\ x^{-1}y^{-2},\ yxy,\ yx^{-1}y\}\\ &\quad\cup\{xy,\ y^{-1}x^{-1},\ xy^{-1},\ yx^{-1},\ y^{-1}x,\ x^{-1}y^{-1}\}\\ &\quad\cup\{y^2,\ y^{-2}\},\end{aligned}$$

has 14 elements, lies in $B(3)$, is non-UP, contains no identity element and has point-group distribution $(n_I, n_X, n_Y, n_Z) = (2, 6, 0, 6)$; its 196 cells fall into 71 coincidence classes, of sizes 10, 8, 6^3, 5^6, 4^5, 2^{55}. In doubled coordinates its elements are

$$\begin{array}{lll}(X,(1,1,0)) & (X,(-1,1,0)) & (X,(1,-1,0))\\ (X,(-1,3,0)) & (X,(-1,1,2)) & (X,(1,1,2))\\ (Z,(1,0,-1)) & (Z,(1,0,1)) & (Z,(1,2,-1))\\ (Z,(1,2,1)) & (Z,(-1,0,1)) & (Z,(-1,2,-1))\\ (I,(0,2,0)) & (I,(0,-2,0)). & \end{array}$$

Proposition 3 (the ball-limited minima)**.** *$B(2)$ contains no non-UP set. For $r = 3$ and $r = 4$, the minimum cardinality of a non-UP subset of $B(r)$ is exactly 14; equivalently, $B(4)$ carries no non-UP set of cardinality at most 13, and A_0 realises 14 inside $B(3)$.*

This is [9, Theorem 3.1] at $r = 2, 3, 4$, reproved here from certificates; the encoding excludes every cardinality from 1 to 13 at once, so the global lower bound $m_1 \geq 8$ of [6] is not used as an input. The radii $r = 5, 6$ of [9] were reproduced here as solver refutations, but their clausal proofs are about 110 and 612 megabytes and were not retained, so they are recorded as measurements and are not theorems here.

The two coordinate functions. For $r \in \{1, 2, 3\}$ set

$$\sigma_r(M, t) = M_{rr} \in \{\pm 1\}, \qquad N_r = \ker(\sigma_r), \qquad \Psi_r(M, t) = t_r \in \mathbb{Z}.$$

Thus $N_1 = \pi^{-1}\{I, X\}$, $N_2 = \pi^{-1}\{I, Y\}$, $N_3 = \pi^{-1}\{I, Z\}$ are the three index-two subgroups of P containing the translation lattice. Two identities carry the whole of §3:

$$\Psi_r(gh) = \Psi_r(g) + \sigma_r(g)\,\Psi_r(h), \qquad \Psi_r(g^{-1}) = -\sigma_r(g)\,\Psi_r(g). \tag{1}$$

Both are immediate from the group law. In particular Ψ_r restricted to N_r is a homomorphism to $\mathbb{Z}$, and its kernel there is exactly the group $L_r \cong \mathbb{Z}^2$ of translations with $t_r = 0$: an element of N_r with $\Psi_r = 0$ has even r-th coordinate and hence, by the parity condition, trivial point part.

3. The peak lemma, the extremal equations, and centring

Throughout this section $A \subseteq P$ is finite and nonempty and $r \in \{1, 2, 3\}$ is fixed; we write $A^0 = A \cap N_r$ and $A^1 = A \setminus N_r$.

Lemma 4 (the located peak lemma)**.** *Let $U, V \subseteq P$ be finite and nonempty, each contained in a single coset of N_r. Then there is a pair $(u, v) \in U \times V$ which maximises $\Psi_r(u'v')$ over $U \times V$ and is the* only *pair of $U \times V$ with product uv.*

Proof sketch. Let ϵ be the common value of σ_r on U. By (1), $\Psi_r(uv) = \Psi_r(u) + \epsilon\Psi_r(v)$, so the maximum of Ψ_r on $U \cdot V$ is attained exactly on $U_{\text{top}} \times V_{\text{ext}}$, where $U_{\text{top}} = \operatorname{argmax} \Psi_r|_U$ and V_{ext} is $\operatorname{argmax} \Psi_r|_V$ or $\operatorname{argmin} \Psi_r|_V$ according to the sign of ϵ. Fix $u_0 \in U_{\text{top}}$ and $v_0 \in V_{\text{ext}}$. By the inverse identity, $u_0^{-1}U_{\text{top}}$ and $V_{\text{ext}}v_0^{-1}$ consist of elements of N_r with $\Psi_r = 0$, hence of translations in L_r. So $U_{\text{top}} \cdot V_{\text{ext}} = u_0\,(S+T)\,v_0$ for finite nonempty sets S, T of translations, and a sumset of translations in a torsion-free abelian group has a uniquely represented element — take the lexicographic maximum in the surviving two coordinates. Translation on either side preserves multiplicities, and every other pair of $U \times V$ has strictly smaller Ψ_r, so that element has multiplicity 1 in $U \cdot V$. ☐

Lemma 4 contains the single-fibre lemma [9, Lemma 4.5], since a single point-group fibre lies inside one coset of N_r for every r. Its existence half is implied by [1], as explained in §1; what the proof adds is that the uniquely represented product is *located* at a Ψ_r-maximum.

Theorem 5 (the extremal equations)**.** *Let $A \subseteq P$ be finite, nonempty and non-UP, and let $r \in \{1,2,3\}$. Then $A^0 \neq \emptyset \neq A^1$, and there is an integer $h_r \geq 0$ with*

$$\textbf{\textit{(E1)}} \quad \max \Psi_r(A^0) = h_r = -\min \Psi_r(A^0),$$
$$\textbf{\textit{(E2)}} \quad \max \Psi_r(A^1) - \min \Psi_r(A^1) = 2h_r.$$

We call $h(A) = (h_1, h_2, h_3)$ the *spread vector* of A; by (E1) it is determined by A.

Proof sketch. If A lay inside one coset of N_r then Lemma 4 applied with $U = V = A$ would give a uniquely represented product, so both parts are nonempty. Because N_r has index 2, the products $A \cdot A$ land in only two cosets of N_r; the coset of the identity receives exactly the two blocks A^0A^0 and A^1A^1, and Ψ_r is a well-defined function on it. If one of those two blocks had a strictly larger Ψ_r-maximum than the other, the element that Lemma 4 produces inside that block would be uniquely represented in the whole of $A \cdot A$: it is unique inside its own block, and the other block does not reach its value. Hence the two maxima are equal, which reads

$$2 \max \Psi_r(A^0) = \max \Psi_r(A^1) - \min \Psi_r(A^1).$$

The same argument run with Ψ_r replaced by $-\Psi_r$ — it is carried out once and instantiated twice — gives

$$-2 \min \Psi_r(A^0) = \max \Psi_r(A^1) - \min \Psi_r(A^1).$$

Comparing the two yields (E1) with $h_r = \max \Psi_r(A^0)$, and either of them is then (E2). Nonnegativity of h_r follows from (E1). ☐

Corollary 6 (three fibres)**.** *A finite nonempty non-UP subset of P meets at least three of the four point-group fibres: for any two classes $c_1, c_2 \in K$ there is $a \in A$ with $\pi(a) \notin \{c_1, c_2\}$.*

Proof. We may assume $c_1 \neq c_2$. One of the three characters $\sigma_1, \sigma_2, \sigma_3$ is constant on $\{c_1, c_2\}$: two distinct diagonal sign matrices of determinant $+1$ agree in exactly one diagonal entry. Say σ_r is constant with value ϵ there. By Theorem 5 the set A has an element with $\sigma_r = -\epsilon$, and that element lies in neither fibre. ☐

This is sharp, and strictly stronger than what [9] proves, namely that A meets at least two fibres: Promislow's A_0 has distribution $(2, 6, 0, 6)$ and so meets exactly three.

Theorem 7 (the centring theorem)**.** *Let $A \subseteq P$ be finite, nonempty and non-UP, with spread vector $h = (h_1, h_2, h_3)$. Then there is a translation c with*

$$cAc^{-1} \ \subseteq\ U(h),$$

where $U(h)$ is the box of §1: the set of $g \in P$ with $|\Psi_r(g)| \le h_r$ for every r with $g \in N_r$, and $|\Psi_r(g)| \le h_r + 1$ for every r with $g \notin N_r$.

Proof. Conjugation by a translation $c = (I, w)$ maps the ambient group to itself and preserves the parity condition, since the shift it induces is even; so it is an automorphism of P, and it preserves non-UP-ness. In doubled coordinates it fixes Ψ_r on N_r and adds the even number $2\Psi_r(c)$ to Ψ_r off N_r, independently in each of the three coordinates. On $A \cap N_r$ the required bound $|\Psi_r| \le h_r$ is (E1), with no conjugation at all. By (E2) the set $\Psi_r(A \setminus N_r)$ lies in a window of width exactly $2h_r$; the target interval $[-(h_r+1), h_r+1]$ has length $2h_r+2$, leaving room 2 for the shift, and every interval of length 2 contains an integer of each parity, so a suitable even shift exists. Choosing w_r for each r independently gives c. □

Because $h(A)$ is defined by (E1) and conjugation by a translation does not move $\Psi_r|_{N_r}$, the spread vector is a conjugation invariant, and Theorem 7 places the conjugate in the box belonging to the original set.

Proposition 8 (from diameter to spread)**.** *Let A be finite, nonempty and non-UP with spread vector h, and suppose* $\mathrm{diam}(A) \le d$*. Then $2h_r \le f_r(d) := \max\{\Psi_r(g) : \ell(g) \le d\}$ for each r.*

Proof. Choose $p, q \in A \cap N_r$ with $\Psi_r(p) = h_r$ and $\Psi_r(q) = -h_r$, which exist by (E1). Since $\sigma_r(q) = +1$, (1) gives $\Psi_r(q^{-1}p) = \Psi_r(p) - \Psi_r(q) = 2h_r$, and $q^{-1}p \in B(d)$ by hypothesis. □

Breadth-first search in the word metric gives $f(1) = (1, 1, 1)$, $f(3) = (3, 3, 2)$, $f(5) = (5, 5, 3)$, $f(7) = (7, 7, 4)$ and $f(9) = (9, 9, 5)$; the values at $d = 3$ and $d = 5$ are used below and are themselves machine-checked, the other three being measurements. Over these values $f_1 = f_2 = d$ and $f_3 = \lceil d/2 \rceil$: the third coordinate is the cheap direction, so the range of h forced by a diameter hypothesis is a flat rectangle rather than a cube: $\mathrm{diam}(A) \le 5$ forces $h \le (2, 2, 1)$, which is 18 profiles with largest box 69 elements, against $h \le (2, 2, 2)$, which is 27 profiles with largest box 99.

4. The box as a decision problem

Fix $h \ge 0$. The box $U(h)$ is a finite subset of P, enumerated by listing, for each of the four point parts M and each coordinate r, the integers of the correct parity class inside the window that M and h_r determine. Its size is at most 99 for all h used here. Note that $U(h)$ can be very degenerate: for instance $U(0, 0, 0) = \{1\}$, because a nontrivial point part forces some t_r to be odd and hence outside a window of width 0.

Proposition 9 (the encoding)**.** *Let U be a finite list of elements of P, let $h \ge 0$ and $k \ge 1$. There is an explicitly constructed propositional formula $\Phi(U, h, k)$ in* *conjunctive normal form with the following property: if $A \subseteq U$ is nonempty, non-UP, has $|A| \le k$ and has spread vector h, then $\Phi(U, h, k)$ is satisfiable. Consequently, if $\Phi(U, h, k)$ is unsatisfiable then no such A exists.*

The formula has one variable sel_i per element of U, one variable $p_{i,j}$ per ordered pair, and the registers of Sinz's sequential counter [8] for the cardinality bound. Its clauses are of six kinds: the two implications defining each $p_{i,j}$; one clause per cell (i, j) saying that some other cell of the $|U| \times |U|$ grid carries the same product; a clause saying at least one element is chosen; the counter clauses; and three families that are theorems of §3 rather than heuristics, namely

- *the fibre cut*: for each pair of point-group classes, A has an element outside their union (Corollary 6);
- *the centring cut*: if $A \cap N_r$ has an element at $\Psi_r = u \ge 0$ then it has one at $\Psi_r \le -u$ (equation (E1));
- *the exactness cut*: h_r and $-h_r$ are both attained on $A \cap N_r$.

The last is the case-split form of (E1), and it is sound only in the presence of the hypothesis "A has spread vector h" in Proposition 9 — which is why every sweep below runs *every* h in its range, so that an arbitrary non-UP set is handled by the instance at its own spread vector. Empirically the exactness cut is what makes the search cheap. Outside the formal development, and recorded here only as a measurement, adding it took the corresponding instance at $h = (2, 2, 2)$ from 20,989,188 bytes of clausal proof to 3,141,647.

Only the direction stated is needed: extra clauses can only make a formula more refutable, so no converse is claimed. Two obligations are therefore discharged for each instance. First, Proposition 9 itself, clause family by clause family. Second, that the formula the solver read is the formula just described: the formula is rendered back to DIMACS inside the proof assistant and compared byte for byte against the file passed to the solver. There are 51 such byte comparisons over the 43 distinct files used below.

Combining Proposition 9 with Theorem 7 gives the pattern used three times in §5: if $A \subseteq P$ is nonempty, non-UP and has spread vector h, then some conjugate $A' = cAc^{-1}$ has the same cardinality, the same spread vector, is still non-UP and lies inside $U(h)$; so an unsatisfiability proof for $\Phi(U(h), h, k)$ excludes A whenever $|A| \le k$.

5. Three sweeps

Theorem 10 (diameter three, any size)**.** *Every finite nonempty non-UP subset $A \subseteq P$ has $h_r(A) \ge 2$ for some $r \in \{1, 2, 3\}$. Consequently no finite nonempty subset of P of word diameter at most 3 is non-UP, whatever its cardinality.*

Proof sketch, and the search that was run. Suppose $h(A) \le (1, 1, 1)$. There are 8 such profiles. For each, the instance $\Phi(U(h), h, k)$ with $k = |U(h)|$ — so that the cardinality bound is vacuous and the counter emits no clauses — was found unsatisfiable, and Proposition 9, applied to the conjugate supplied by Theorem 7, then contradicts the existence of A. The box sizes are $|U(0, 0, 0)| = 1$, $|U(h)| = 5$ for the three profiles with one coordinate equal to 1, 17 for the three with two, and $|U(1, 1, 1)| = 37$. Four of the eight profiles are refuted inside the proof kernel, with no solver involved: their boxes meet only two point-group fibres, so the corresponding fibre cut is a clause with no literals at all. The other four are refutations found by a SAT solver and replayed from clausal proofs totalling 101,427 bytes, the largest being 80,472 bytes at $h = (1, 1, 1)$.

For the diameter form, $\mathrm{diam}(A) \leq 3$ forces $2h_1 \leq 3$, $2h_2 \leq 3$, $2h_3 \leq 2$ by Proposition 8 and $f(3) = (3, 3, 2)$, hence $h \leq (1, 1, 1)$. □

Theorem 11 (diameter five, size at most thirteen)**.** *No finite nonempty non-UP subset of P of cardinality at most 13 has word diameter at most 5. Equivalently, every non-UP subset of P of word diameter at most 5 has at least 14 elements.*

Proof sketch, and the search that was run. By Proposition 8 and $f(5) = (5, 5, 3)$, a word diameter of at most 5 forces $h \leq (2, 2, 1)$, which is 18 profiles. Each instance $\Phi(U(h), h, 13)$ is unsatisfiable. Six of the eighteen are refuted in the kernel by the empty fibre-cut clause; the other twelve are replayed from clausal proofs totalling 1,334,304 bytes, the largest of them 480,200 bytes at $h = (2, 2, 1)$, where the box has 69 elements. The conjugate supplied by Theorem 7 has the same cardinality as A. □

For comparison, the certified statement about $B(4)$ in Proposition 3 costs about 2.3 megabytes of certificate and speaks about 83 group elements; the sweep above costs 1.33 megabytes and speaks about the whole group.

Theorem 12 (classification of the small spread profiles)**.** *Let $A \subseteq P$ be finite, nonempty and non-UP, with spread vector h. Then either $h_r \geq 3$ for some r, or h is one of $(1, 1, 2)$, $(1, 2, 1)$, $(2, 1, 1)$ and $|A| \geq 14$. In particular, if $h_r \leq 2$ for all r then h is a cyclic rotation of Promislow's profile $(1, 2, 1)$ and A has at least 14 elements.*

Proof sketch, and the search that was run. There are 27 profiles with $h \leq (2, 2, 2)$. For each, the instance $\Phi(U(h), h, k)$ with $k = |U(h)|$ was run. Exactly 24 are unsatisfiable: seven of these are refuted in the kernel by an empty fibre-cut clause, and seventeen are replayed from clausal proofs totalling 6,850,190 bytes, the largest being 3,762,348 bytes at $h = (2, 2, 2)$, where the box has 99 elements. The remaining three profiles $(1, 1, 2)$, $(1, 2, 1)$, $(2, 1, 1)$ are satisfiable, so no refutation exists for them and none is claimed; for the size half, the instances $\Phi(U(h), h, 13)$ at those three profiles were run and are unsatisfiable. Two of the three, at $(1, 2, 1)$ and $(2, 1, 1)$, are already part of the rectangle of Theorem 11; the third, at $(1, 1, 2)$, has a box of 51 elements and a clausal proof of 588,497 bytes. Every box here has at most 99 elements, so the whole classification is decided over universes substantially smaller than the balls $B(6)$ and $B(7)$ that [9] searches. □

Remark 13. The three satisfiable profiles form a single orbit under the cyclic rotation of the three coordinates. That rotation is an automorphism of P: it is the 3-cycle $\sigma\colon (a, b, c) \mapsto (c, a, b)$ of [9, Remark 2.2], where the automorphism group of P is recorded as containing a copy of S_3 acting by coordinate permutations, and where it is also observed that σ preserves no ball about the identity, since $\sigma(y)$ is not a generator. That is a plausible reason why a ball search of small radius sees only the orbit of the generator swap $x \leftrightarrow y$ — [9, Proposition 4.3] finds exactly two point-group distributions of a non-UP 14-set, inside $B(3)$ and inside $B(4)$, and they are swapped by it — while the classification above sees a full 3-cycle orbit of spread profiles. The box $U(2, 2, 2)$ has 99 elements and $U(2, 2, 1)$ has 69; the smallest ball containing either is $B(7)$, not $B(6)$, which is a measurement made outside the formal development. If it is right, both sweeps already look at group elements outside the largest ball for which [9] reports a certified computation.

6. Sharpness, controls, and what the sweeps do not say

Proposition 14 (controls)**.** *Let A_0 be Promislow's 14-element set of Proposition 2. Then*

(i) *the hypothesis of Theorem 10 is satisfiable: $B(1)$, a 5-element set, has word diameter at most 3;*
(ii) *A_0 has word diameter at most 6, and does* not *have word diameter at most 5;*
(iii) *the spread vector of A_0 is $(1, 2, 1)$, computed through the same definitions used in §3;*
(iv) *$A_0 \subseteq U(1, 2, 1)$, with no conjugation applied.*

Item (ii) is the negative control for the diameter statements: since A_0 is non-UP, the bound 5 of Theorem 11 cannot be raised to 6, and the unconditional bound 3 of Theorem 10 cannot be raised past 5. Item (iii) makes Theorem 10 sharp, $\max_r h_r = 2$ being attained, and puts Promislow's set at one of the three profiles of Theorem 12. Item (iv) is the load-bearing calibration of the encoding: it forces the unbounded instance at $h = (1, 2, 1)$ to be satisfiable, and it is, so none of the three added clause families is over-constraining. Had any of them been, that instance would have come out unsatisfiable and the classification would have listed four or more refuted profiles. Independently, seven profiles are refuted twice — once in the kernel by the empty fibre clause and once by the solver — by routes that share no code below the formula. Finally, (E1) is not a tautology. A translation $c = (I, w)$ lies in every N_r, so $cA_0 \cap N_r = c\,(A_0 \cap N_r)$, and by (1) the values of Ψ_r there are those on $A_0 \cap N_r$ shifted by w_r; hence (E1) fails for cA_0 in every coordinate with $w_r \neq 0$, and $w \neq 0$ for every nontrivial translation. That is the same phenomenon as the failure of translation invariance recorded in [9, Remark 2.3], where every nontrivial translate of A_0 inside $B(3)$ is found to acquire a unique product.

Remark 15 (what is not proved). Theorem 11 does not bound the diameter of a minimal non-UP set, and neither does anything else here. All three theorems have the form "no non-UP set with small diameter or small spread", never "every non-UP set has small diameter". In particular $m_1(P) \in [8, 14]$ and [9, Question 3.2] are exactly where [9] left them. What the centring theorem changes is that the question becomes one about a single integer invariant with a cheap decision procedure at each value, rather than about the location of a hypothetical witness.

Proposition 16 (the two-sided minimum inside a ball)**.** *The minimum of $|A|+|B|$ over non-UP pairs (A, B) with $A, B \subseteq B(3)$ is exactly 24, and no non-UP pair has both sides inside $B(2)$.*

This is [9, Proposition 3.3], both halves, reproved here from certificates.

Remark 17 (the two-sided problem). For pairs, bi-translation is a genuine symmetry, and it uses up exactly the freedom that (E1) pins down in the one-set case; there is no absolute centring for pairs, and the box method of §4 has no pair analogue at present. The published interval $m_2(P) \in [16, 24]$ therefore stands: its upper end is the pair of Proposition 16 and its lower end is [6]. Outside the formal development, and reported here only as a computation, a solver refutation whose clausal proof is

8.3 gigabytes — too large to retain or to check — shows in addition that no non-UP pair of total size at most 23 has $A \subseteq a_0B(4)$ for some $a_0 \in A$ and $B \subseteq B(4)b_0$ for some $b_0 \in B$. That normalisation is the one [9] itself uses at radius 4 for totals up to 19; the computation extends the range of totals and nothing else, and it does not move either end of $[16, 24]$.

Remark 18 (cost, and the next range). The three sweeps together are 46 solver instances over boxes of at most 99 elements and 43 refutations. The next diameter, 7, forces $h \leq (3, 3, 2)$: 48 profiles with largest box 195 elements. It looks reachable: 46 of the 48 instances have been run and are unsatisfiable. But at this decomposition three of them individually produce clausal proofs of 19 to 32 megabytes, which is why diameter 7 is not claimed here. Both statements in this remark are measurements, made outside the formal development.

7. Verification

Every theorem and proposition above is machine-checked in Lean 4 [5] against Mathlib. The group is carried as an exact integer model — a point part stored as two bits, the translation part doubled — so that no rational number appears and the group law, the parity condition, the generators and their relations, torsion-freeness and the word-metric ball are ordinary definitions with ordinary proofs. The whole of §3 is proved without any computation: the two identities (1) hold by definitional unfolding in each of the three coordinates, and the peak lemma, the extremal equations, the three-fibre corollary, the centring theorem and the diameter conversion are proved from them; the only axioms in the centring theorem's proof are propositional extensionality, quotient soundness and choice, with no evaluation of any kind. The finite content is confined to §§4–5: the enumeration of each box, the rendering of each formula, the byte comparison against the DIMACS file the solver read, and the replay of each clausal proof are all carried out by compiled evaluation inside Lean and so rest additionally on Lean's native-evaluation axiom, as do the two breadth-first tables $f(3)$ and $f(5)$ and the controls of Proposition 14 other than (iii), which is a kernel computation. Refutations were produced with CaDiCaL 2.1.2 [2] emitting LRAT proofs [3], and are replayed through Lean's verified LRAT checker; 13 of the 43 refutations need no certificate at all, being instances in which one clause of the fibre cut is empty, and are discharged by the kernel. The certificates retained total 8,772,991 bytes in 30 files, against 2.3 megabytes for the single ball statement at $r = 4$. No proof uses `sorry`, and the two generators that produce the DIMACS files ship with self-tests reproducing, respectively, the published ball sizes, relations and witness, and the three sweep totals. Repository reference to be added at submission.

References

[1] B. H. Bowditch, *A variation on the unique product property*, J. London Math. Soc. (2) **62** (2000), no. 3, 813–826, doi:10.1112/S0024610700001307.

[2] A. Biere, T. Faller, K. Fazekas, M. Fleury, N. Froleyks and F. Pollitt, *CaDiCaL 2.0*, in: Computer Aided Verification (CAV 2024), Lecture Notes in Computer Science **14681**, Springer, 2024, pp. 133–152. The refutations reported here were run with CaDiCaL 2.1.2, the build shipped with the Lean toolchain used for the formal development.

[3] L. Cruz-Filipe, M. J. H. Heule, W. A. Hunt Jr., M. Kaufmann and P. Schneider-Kamp, *Efficient certified RAT verification*, in: Automated Deduction — CADE 26, Lecture Notes in Computer Science **10395**, Springer, 2017, pp. 220–236.

[4] G. Gardam, *A counterexample to the unit conjecture for group rings*, Ann. of Math. (2) **194** (2021), 967–979.

[5] 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, pp. 625–635.

[6] P. P. Nielsen and L. Soelberg, *Small sets without unique products in torsion-free groups*, J. Algebra Appl. **23** (2024), no. 8, Paper No. 2550050, doi:10.1142/S0219498825500501. We did not read this paper; it is cited for the bounds $m_1 \geq 8$ and $m_2 \geq 16$ as reported in [9].

[7] S. D. Promislow, *A simple example of a torsion-free, non unique product group*, Bull. London Math. Soc. **20** (1988), no. 4, 302–304, doi:10.1112/blms/20.4.302.

[8] C. Sinz, *Towards an optimal CNF encoding of Boolean cardinality constraints*, in: Principles and Practice of Constraint Programming — CP 2005, Lecture Notes in Computer Science **3709**, Springer, 2005, pp. 827–831.

[9] M. Tabei, *Least sizes of non-unique-product sets: the Promislow group and a Heisenberg-type candidate*, arXiv:2607.18346v1, 20 July 2026. Section, theorem, proposition, remark and table numbers cited here are those of v1, which is the only version, and carries no journal reference, as of 30 August 2026. A companion preprint, M. Tabei, *The quantitative non-unique-product landscape at the global minimum: the Nielsen–Soelberg groups*, arXiv:2607.19687, 22 July 2026, treats other groups and repeats the scope note that symmetric non-UP is not translation invariant, so that no diameter statement follows from a ball computation.