

The Polar-Code Endpoint for Class-Two p -Groups

Luca Blanchi

June 2026

Abstract

Let p be odd and let G_β be a special class-two exponent- p group with commutator tensor

$$\beta : \Lambda_{\mathbb{F}_p}^2 V \rightarrow W.$$

Writing

$$C_\beta = \beta^*(W^*) \leq \Lambda^2 V^*,$$

isomorphism of the groups is equivalent to projective equivalence of the alternating-form system C_β . This capstone note records the endpoint of the observable hierarchy developed in the surrounding class-two group papers. Ordinary word-map distributions factor through the scalar rank profile

$$[\omega] \mapsto \text{rank } \omega, \quad [\omega] \in \mathbb{P}(C_\beta),$$

and therefore leave large spectral and Pfaffian fibres. At the other end, the polar code

$$\text{Pol}(C) = \{([\omega], [u], [v]) \in \mathbb{P}(C) \times \mathbb{P}(V) \times \mathbb{P}(V) : \omega(u, v) = 0\}$$

is complete: it determines C up to the natural $\text{PGL}(V)$ -action, and hence determines G_β up to isomorphism.

The main contribution is to make this endpoint usable without scattering the new material across the earlier papers. We prove the polar-code completeness theorem, an unconditional frame normal form, a landmark-frame normal form with cost $q^{O(nb)}$, and an elementary generic stabilizer estimate for the action of $\text{PGL}(V)$ on $\text{Gr}(m, \Lambda^2 V^*)$. We then isolate two geometric compression modules: projected-Grassmannian reconstruction under a secant-avoidance hypothesis, and slice-gluing under explicit overlap-rigidity hypotheses. These modules explain how quadratic word observables, support-image profiles, exact pencil thresholds, commutator-word tomography of pencils, odd-Pfaffian reconstruction, and Pfaffian word-measure twins fit into a single fibre-resolution ladder.

The final sections preserve the stronger structural ideas suggested by the development, but with their mathematical status separated. Projection-stable and slice-rigid loci give cap-free generic reconstruction. The remaining bad locus is organized by low-rank, persistent Pfaffian singular, adjoint-decomposable, field-extension, tensor-product, classical-polar, and primitive polar-core behaviours. For primitive cores we state a relative standard-structure reduction: assuming recognizers and canonizers for the usual geometric branches, Aschbacher-type structure and bounded-base results reduce the residual problem to standard geometry or small-base canonicalization. Thus the paper does not claim a polynomial-time worst-case algorithm for alternating matrix-space isometry. It identifies the complete invariant, the generic geometric compression mechanisms, and the precise recognizer/canonizer bottleneck left by the worst case.

Declaration of generative AI and AI-assisted technologies in the writing process

This article belongs to the Presentation theory section. Generative AI, including ChatGPT by OpenAI, was used to assist with mathematical drafting, formalization, review, and editing. The

note may be preliminary, may have been only partially reviewed, and in some cases may be only AI-reviewed.

1 Introduction

Special finite p -groups of nilpotency class two and exponent p , with p odd, are controlled by alternating bilinear algebra. If

$$V = G/Z(G), \quad W = G' = Z(G),$$

then the commutator gives an alternating tensor

$$\beta : \Lambda^2 V \rightarrow W.$$

Conversely, such a tensor defines a group on $V \oplus W$ by

$$(v, z)(v', z') = \left(v + v', z + z' + \frac{1}{2}\beta(v, v') \right),$$

with commutator

$$[(v, z), (v', z')] = (0, \beta(v, v')).$$

The group is special precisely when the tensor is surjective and has zero common radical.

The isomorphism problem for these groups is therefore the pseudo-isometry problem for alternating tensors. Equivalently, it is the problem of classifying the m -dimensional subspace

$$C_\beta = \beta^*(W^*) \leq \Lambda^2 V^*, \quad m = \dim W,$$

under change of basis in V . This is the alternating matrix-space isometry problem in geometric language.

The articles in this cluster study several observables of C_β . Ordinary word maps see only scalar rank data. Support-image profiles add codomain-localized rank information. Fourier-tagged commutator-word moments can recover full Kronecker data for pencils and, on generic odd-Pfaffian loci, higher-dimensional Pfaffian data. Kernel incidence resolves Pfaffian fibres invisible to ordinary word measures. The polar code introduced below is the endpoint of this hierarchy: it is the full incidence relation of the polarities defined by the forms in C_β .

This note is the endpoint of that sequence, not a replacement for the specialized papers. Its role is to collect the global material that would otherwise be awkwardly duplicated: complete polar-code observability, canonical normal forms, generic geometric reconstruction, and the residual standard-structure reduction.

2 Class-two groups and alternating systems

Let $k = \mathbb{F}_q$ have odd characteristic. For a finite-dimensional k -vector space V , write

$$R = \Lambda^2 V^*, \quad r = \dim R = \binom{\dim V}{2}.$$

An alternating system is a subspace

$$C \leq R.$$

It defines a target-dual tensor

$$\beta_C : \Lambda^2 V \rightarrow C^*$$

by evaluation. When β_C is surjective and radical-free, it defines a special class-two exponent- p group.

Proposition 2.1 (Group isomorphism and alternating systems). *Let G_β and $G_{\beta'}$ be special class-two exponent- p groups with commutator tensors*

$$\beta : \Lambda^2 V \rightarrow W, \quad \beta' : \Lambda^2 V' \rightarrow W'.$$

Put

$$C_\beta = \beta^*(W^*) \leq \Lambda^2 V^*, \quad C_{\beta'} = \beta'^*(W'^*) \leq \Lambda^2 V'^*.$$

Then $G_\beta \cong G_{\beta'}$ if and only if there is a linear isomorphism $g : V \rightarrow V'$ such that

$$(\Lambda^2 g^*)C_{\beta'} = C_\beta.$$

Proof. An isomorphism of special groups induces linear isomorphisms

$$V \rightarrow V', \quad W \rightarrow W',$$

compatible with commutators. Thus the commutator tensors are pseudo-isometric. Dualizing the target change identifies this with equality of the pulled-back systems of scalar alternating forms after the induced change of basis on V .

Conversely, an equivalence of alternating systems gives a pseudo-isometry of the tensors after choosing the induced target map $W \rightarrow W'$. The standard Baer correspondence between class-two exponent- p groups and alternating tensors then gives an isomorphism of the associated groups. $\square$

3 Word observables and scalar rank

The weakest natural observables in this setting are ordinary word-map distributions. They are important because they are intrinsic group-theoretic data, but in class two they collapse to scalar rank.

Definition 3.1 (Scalar rank profile). *For an alternating system $C \leq \Lambda^2 V^*$, its scalar rank profile is the function*

$$R_C : \mathbb{P}(C) \rightarrow \mathbb{Z}_{\geq 0}, \quad R_C([\omega]) = \text{rank } \omega.$$

Proposition 3.2 (Ordinary word maps factor through scalar rank). *For class-two exponent- p groups G_β , the collection of all ordinary word-map distributions is determined by*

$$\dim V, \quad \dim W, \quad R_{C_\beta}.$$

In particular, two tensors with the same scalar rank profile have the same ordinary word-map distributions.

Proof. In the free class-two exponent- p group on s variables, every word has a unique normal form consisting of a linear exponent-sum part and an alternating commutator part:

$$w(x_1, \dots, x_s) = \sum_i a_i x_i + \sum_{i < j} b_{ij} [x_i, x_j].$$

Evaluating on $G_\beta = V \oplus W$, the central component is a finite linear combination of scalar alternating forms $\lambda \circ \beta$ after Fourier expansion in the target W . For each Fourier frequency, the number of solutions is governed by the rank of the corresponding scalar alternating form. Summing the Fourier terms gives a quantity depending only on $\dim V$, $\dim W$, and the rank profile R_{C_β} . $\square$

Thus ordinary word measures are not complete. They do not see the full projective polarity associated with each scalar alternating form; they see only its rank.

4 The polar code

The missing data are the polarities themselves.

Definition 4.1 (Polar code). *Let $C \leq \Lambda^2 V^*$ be a nonzero alternating system. Its polar code is the incidence variety*

$$\text{Pol}(C) = \{([\omega], [u], [v]) \in \mathbb{P}(C) \times \mathbb{P}(V) \times \mathbb{P}(V) : \omega(u, v) = 0\}.$$

An isomorphism of polar codes from C to C' is a pair of projective isomorphisms

$$\phi : \mathbb{P}(C) \rightarrow \mathbb{P}(C'), \quad \psi : \mathbb{P}(V) \rightarrow \mathbb{P}(V')$$

such that

$$([\omega], [u], [v]) \in \text{Pol}(C) \iff (\phi[\omega], \psi[u], \psi[v]) \in \text{Pol}(C').$$

Lemma 4.2 (A projective polarity determines the alternating form). *Assume $\dim V \geq 3$. Let $\omega, \eta \in \Lambda^2 V^*$ be nonzero alternating forms. If*

$$\omega(u, v) = 0 \iff \eta(u, v) = 0$$

for all projective points $[u], [v] \in \mathbb{P}(V)$, then $\eta = c\omega$ for some $c \in k^\times$.

Proof. For each nonzero $u \in V$, the relation determines the hyperplane

$$u_\omega^\perp = \{v \in V : \omega(u, v) = 0\}.$$

Thus the two linear functionals

$$\omega(u, -), \quad \eta(u, -)$$

have the same kernel. If $\omega(u, -) \neq 0$, then

$$\eta(u, -) = c(u)\omega(u, -)$$

for some scalar $c(u)$. Choose u_1, u_2 such that the two functionals $\omega(u_1, -)$ and $\omega(u_2, -)$ are nonzero and not proportional. Applying the same comparison to $u_1 + u_2$ gives

$$c(u_1 + u_2)(\omega(u_1, -) + \omega(u_2, -)) = c(u_1)\omega(u_1, -) + c(u_2)\omega(u_2, -),$$

so $c(u_1) = c(u_2) = c(u_1 + u_2)$. By moving through pairs whose polar functionals are not proportional, the scalar is constant on the complement of the radical. On the radical both forms vanish. Hence $\eta = c\omega$. $\square$

Theorem 4.3 (Completeness of the polar code). *Let $\dim V \geq 3$. The polar code $\text{Pol}(C)$ determines C up to the natural projective action of $\text{PGL}(V)$. Consequently, for special class-two exponent- p groups,*

$$\text{Pol}(C_\beta) \cong \text{Pol}(C_{\beta'}) \iff G_\beta \cong G_{\beta'}.$$

Proof. For each point $[\omega] \in \mathbb{P}(C)$, the fibre of the polar code over $[\omega]$ is the projective orthogonality relation of the alternating form ω . By the preceding lemma, this fibre recovers the point

$$[\omega] \in \mathbb{P}(\Lambda^2 V^*).$$

As $[\omega]$ varies in $\mathbb{P}(C)$, the polar code therefore recovers the embedded projective subspace

$$\mathbb{P}(C) \subset \mathbb{P}(\Lambda^2 V^*).$$

Thus it recovers C up to the projective change of coordinates induced by the underlying isomorphism of $\mathbb{P}(V)$. The group statement follows from Proposition 2.1. $\square$

This theorem is the exact endpoint of the observable hierarchy. The scalar rank profile remembers only the rank of each polarity. The polar code remembers the polarity itself.

5 Normal forms

Completeness gives an invariant. A normal form gives a canonical representative.

Definition 5.1 (Frame normal form). *Fix ordered bases of V and C . Represent the subspace*

$$C \leq \Lambda^2 V^*$$

by the reduced row echelon form of a matrix of Plucker coordinates in $\Lambda^2 V^$. Define*

$$\text{NF}_{\text{frame}}(C) = \min_{\text{lex}} \left\{ \text{RREF}((\Lambda^2 g)C) : g \in \text{PGL}(V) \right\}.$$

Over a finite field this minimum is taken over a finite set.

Theorem 5.2 (Unconditional frame normal form). *For $C, C' \leq \Lambda^2 V^*$ of the same dimension,*

$$C \sim C' \iff \text{NF}_{\text{frame}}(C) = \text{NF}_{\text{frame}}(C').$$

For special class-two exponent- p groups this gives a complete isomorphism normal form.

Proof. If $C' = hC$, then the two $\text{PGL}(V)$ -orbits give the same set of reduced row echelon representatives, hence the same lexicographic minimum. Conversely, if the minima agree, some representatives in the two orbits are equal, so C and C' are projectively equivalent. Proposition 2.1 transfers the statement to groups. $\square$

The frame normal form is complete but not compressed. The presentation-theoretic question is therefore not whether a complete invariant exists; it does. The question is which intrinsic data reduce the amount of orbit enumeration.

Definition 5.3 (Polar landmark frame). *A polar landmark frame for C is a finite set*

$$B \subset \mathbb{P}(V) \cup \mathbb{P}(C)$$

with the following properties.

- i. B is defined intrinsically from the polar code by incidence and rank data.*
- ii. From B one constructs a projective frame of $\mathbb{P}(V)$.*

iii. The pointwise stabilizer of B in $\text{Aut}(\text{Pol}(C))$ is trivial.

Let $b_{\text{pol}}(C)$ be the least size of such a landmark frame, when one exists.

Theorem 5.4 (Canonicalization from landmarks). *Suppose $C \leq \Lambda^2 V^*$ over $\mathbb{F}_q$, with $\dim V = n$, has a polar landmark frame of size b . Then C has a canonical normal form obtained by enumerating at most*

$$q^{O(nb)}$$

landmark choices, followed by polynomial-time linear algebra in the ambient coordinates.

Proof. For each candidate landmark set B , construct the associated projective frame F_B and apply the unique element of $\text{PGL}(V)$ carrying F_B to the standard frame. Compute the reduced row echelon form of the transformed subspace. Because the landmark condition is intrinsic, isomorphic polar codes produce corresponding candidate sets. Because the pointwise stabilizer is trivial, no residual projective ambiguity remains after the frame is fixed. Taking the lexicographic minimum over all intrinsic landmark choices gives a canonical representative.

The number of possible projective points in $\mathbb{P}(V)$ is $O(q^{n-1})$, and the number of projective points in $\mathbb{P}(C)$ is at most $O(q^{m-1})$ with $m \leq \binom{n}{2}$. Thus the number of b -tuples to enumerate is $q^{O(nb)}$ for the parameter ranges considered here. The remaining steps are row reductions and comparisons. $\square$

6 Generic rigidity

The previous theorem is useful when small landmarks exist. A first general source of compression is generic triviality of stabilizers.

Lemma 6.1 (Fixed-subspace codimension). *Let R be an r -dimensional vector space over a field, and let $T \in \text{GL}(R)$ be non-scalar. For $1 \leq m \leq r-1$, the locus of m -dimensional T -stable subspaces in $\text{Gr}(m, R)$ has codimension at least*

$$\min(m, r - m).$$

Proof. After base change to an algebraic closure, the fixed locus is a union of varieties of submodules for the $k[T]$ -module R . The largest possible fixed locus for a non-scalar operator occurs in the coarsest degeneration: either a decomposition into an $(r-1)$ -dimensional block and a line, or the corresponding rank-one unipotent limit. In that case a stable m -plane either lies in the hyperplane, contains the line, or is compatible with the two-step flag. The largest component has dimension

$$m(r - m) - \min(m, r - m).$$

Refining the invariant decomposition or adding nontrivial nilpotent conditions only imposes further incidence conditions. Thus the codimension is at least $\min(m, r - m)$. $\square$

Theorem 6.2 (Generic stabilizer estimate). *Let $n \geq 3$, let*

$$R = \Lambda^2 V^*, \quad r = \binom{n}{2},$$

and let $\text{PGL}(V)$ act on $\text{Gr}(m, R)$. Put

$$d = \min(m, r - m).$$

If $d > n^2 + 1$, then a uniformly random m -subspace

$$C \leq R$$

over $\mathbb{F}_q$ has trivial stabilizer in $\mathrm{PGL}(V)$ with probability

$$1 - O(q^{-1})$$

as $q \rightarrow \infty$ with n and m fixed.

Proof. For every nontrivial $g \in \mathrm{PGL}(V)$, the induced action on $R = \Lambda^2 V^*$ is non-scalar when $n \geq 3$. By the codimension lemma, the locus of m -subspaces fixed by g has codimension at least d in $\mathrm{Gr}(m, R)$. Hence the probability that a random C is fixed by this g is $O(q^{-d})$.

The group $\mathrm{PGL}(V)(\mathbb{F}_q)$ has $O(q^{n^2-1})$ elements. A union bound gives

$$\Pr[\exists g \neq 1 \text{ with } gC = C] \leq O(q^{n^2-1-d}).$$

If $d > n^2 + 1$, this is $O(q^{-1})$ or better. Thus the stabilizer is generically trivial in the stated range. $\square$

Remark 6.3. *This theorem is deliberately modest. It is a finite-field generic estimate, not a classification of all stabilizers and not a worst-case algorithm. Its role is to justify why polar-code refinement should often find small intrinsic frames in broad middle-dimensional ranges.*

7 Two reconstruction modules

The next two statements are safe modules. They are useful because their hypotheses are explicit. They should not be read as unconditional generic classification theorems in every range.

7.1 Projected Grassmannians

Let

$$C \leq \Lambda^2 V^*, \quad K = C^\perp \leq \Lambda^2 V.$$

Projection from the center $\mathbb{P}(K)$ gives a rational map

$$\pi_C : \mathbb{P}(\Lambda^2 V) \dashrightarrow \mathbb{P}(C^*).$$

Restrict it to the Plucker Grassmannian

$$\mathrm{Gr}_2(V) \subset \mathbb{P}(\Lambda^2 V),$$

and denote the image by

$$X_C = \pi_C(\mathrm{Gr}_2(V)).$$

Proposition 7.1 (Secant-avoidance reconstruction template). *Assume $\mathbb{P}(K)$ is disjoint from the secant variety*

$$\sigma_2(\mathrm{Gr}_2(V)).$$

Then

$$\mathrm{Gr}_2(V) \rightarrow X_C$$

is an embedding. Moreover, if the embedded variety $X_C \subset \mathbb{P}(C^)$ is recovered together with its hyperplane coordinates, then C is recovered up to the natural projective action on V , up to the standard Grassmannian duality ambiguity in the exceptional self-dual cases.*

Proof. Avoiding the secant variety means that no secant or tangent direction of the Grassmannian is collapsed by the projection. Therefore the restriction of the projection to $\text{Gr}_2(V)$ is an embedding.

The hyperplane coordinates of the embedded image are precisely the linear subsystem $C \subset H^0(\text{Gr}_2(V), \mathcal{O}(1)) = \Lambda^2 V^*$. Projective automorphisms of the Plucker Grassmannian are induced by semilinear transformations of V , with the usual duality exception. Hence the embedded projected Grassmannian recovers the subsystem C up to the corresponding projective change of coordinates. $\square$

The dimension of the rank- ≤ 4 skew locus is

$$\dim \sigma_2(\text{Gr}_2(V)) = 4n - 11.$$

Thus a generic center avoids it whenever

$$m \geq 4n - 10.$$

The dual inequality

$$\binom{n}{2} - m \geq 4n - 10$$

gives the analogous complementary range after passing to annihilators.

7.2 Slice gluing

The middle range often requires reconstruction from smaller subsystems.

Definition 7.2 (Rigid slice cover). *A rigid slice cover of C is a finite family of subspaces*

$$U_i \leq C$$

such that:

- i. the U_i generate C ;*
- ii. each U_i is reconstructible up to projective equivalence;*
- iii. the graph with vertices i and edges $i \sim j$ when $U_i \cap U_j$ is reconstructibly nontrivial is connected;*
- iv. every overlap used for gluing has trivial residual ambiguity inside the reconstructed slice data;*
- v. the global residual stabilizer after gluing is scalar.*

Proposition 7.3 (Slice-gluing reconstruction). *If C admits a rigid slice cover, then C is reconstructible up to projective equivalence.*

Proof. Reconstruct the first slice. Along an edge of the overlap graph, the common reconstructed overlap identifies the next slice in the same projective gauge, because the overlap has no residual ambiguity. Since the overlap graph is connected, all slices are glued in one projective coordinate system. The slices generate C , so their union determines C . The final scalar stabilizer condition removes the last projective ambiguity. $\square$

In applications, 2-dimensional slices are alternating pencils and can be handled by Kronecker data; 3-dimensional slices lead to Pfaffian curves and kernel incidence; odd-dimensional Pfaffian loci give another reconstruction mechanism on suitable open sets. The slice-gluing proposition is the formal transport step that lets those local reconstructions become a global reconstruction when the overlaps are rigid.

8 Generic cap-free reconstruction

The previous two modules give a clean generic statement. The point is not that all fibres are now easy. It is that the brute-force separating cap is not part of the generic geometry.

Definition 8.1 (Good slice locus). *For $s \geq 2$, let*

$$\mathcal{G}_s \subset \mathrm{Gr}(s, \Lambda^2 V^*)$$

be the locus of s -dimensional slices satisfying the local hypotheses needed by the chosen reconstruction method: radical-freeness where appropriate, expected rank behaviour, smooth or geometrically reduced degeneracy data, scalar common self-adjoint algebra when used, and reconstructibility from the local observable assigned to that slice.

For $s = 2$, this is the regular pencil locus. For $s = 3$, it is the Pfaffian-net locus where the curve and kernel incidence behave well. For $s = 4$ in odd dimension, it is the odd-Pfaffian locus where the Buchsbaum–Eisenbud reconstruction mechanism applies. In each case the conditions are open. The non-vacuity of the relevant open locus is part of the local input.

Theorem 8.2 (Cap-free generic reconstruction on controlled loci). *Let $C \leq \Lambda^2 V^*$ be an m -dimensional alternating system over $\mathbb{F}_q$, with $n = \dim V$ and $r = \binom{n}{2}$.*

- i. If $m \leq 2$, then the corresponding special class-two exponent- p groups are handled by the one-form case and by Kronecker/Fourier-tagged pencil tomography.*
- ii. If $m \geq 4n - 10$ or $r - m \geq 4n - 10$, then a generic C is reconstructed from the projected Grassmannian module.*
- iii. If C admits a rigid cover by good slices, then C is reconstructed by slice gluing.*

In these cases the polar-code endpoint is reached without an orbit-indicator separating cap.

Proof. The case $m = 1$ is the classification of one alternating form by rank. The case $m = 2$ is the alternating-pencil case, where Kronecker data determine the pencil up to projective target change and congruence of the source.

In the projection-stable range, Proposition 7.1 applies: the projection center avoids the secant variety of $\mathrm{Gr}_2(V)$ generically, so the projected Grassmannian is embedded and recovers the subsystem C up to the natural projective action.

In the middle range, the assertion is exactly Proposition 7.3 once the good slices and rigid overlaps are supplied. The slices reconstruct in one projective gauge and generate C . $\square$

Remark 8.3. *This theorem deliberately distinguishes a proved gluing mechanism from the stronger expectation that every generic middle-range system admits such a cover. The latter is a natural generic slice-rigidity problem. It follows on any nonempty open locus where the good-slice and overlap-rigidity hypotheses can be verified.*

9 Residual geometry

The complement of the controlled loci is not discarded. It is the residual geometry of the problem. A useful diagnostic definition is

$$\mathcal{S}_{n,m}^{\mathrm{bad}} = \{C : C \text{ is not projection-stable, not slice-rigid, and not adjoint-decomposable}\}.$$

This is a shorthand for a union of closed or locally closed failure conditions: the projection center meets a low-rank secant stratum; the chosen slice loci are singular or nonfaithful; overlaps have residual ambiguity; or the common adjoint algebra does not expose a smaller module structure.

Definition 9.1 (Residual types). *The residual behaviours relevant to the polar-code endpoint are:*

1. *low-rank or secant-intersection type;*
2. *persistent singular Pfaffian type;*
3. *non-scalar adjoint type;*
4. *field-extension type;*
5. *tensor-product or imprimitive type;*
6. *classical-polar type;*
7. *primitive polar-core type.*

The list should be read as a structural atlas rather than as a fully proved irreducible-component classification in this note. Each item has a concrete meaning.

Low-rank/secant. Here $C^\perp$ meets the rank- ≤ 4 secant of $\text{Gr}_2(V)$. The support variety

$$S(C) = \{U \in \text{Gr}_4(V) : \mathbb{P}(\Lambda^2 U) \cap \mathbb{P}(C^\perp) \neq \emptyset\}$$

records the low-rank obstruction.

Persistent singular Pfaffian. The rank loci

$$Z_C^{(a)} = \{[\omega] \in \mathbb{P}(C) : \text{rank } \omega \leq a\}$$

are singular, nonreduced, nonradical, or have kernel incidence that is not faithful over a dominant family of slices.

Non-scalar adjoint. The common self-adjoint algebra is larger than the scalar algebra. Its idempotents, radical, semisimple quotient, or field subalgebras produce module decompositions, field-extension structures, or tensor structures.

Standard geometric structure. Field-extension, tensor-product, imprimitive, subfield, tensor-induced, symplectic-type, and classical-form behaviours are standard linear structures on V preserved by the residual automorphism group of the polar code.

Primitive polar core. No smaller structure has yet been exposed. The polar code is still complete, and frame normalization still gives a finite canonical normal form, but efficient normalization has not been reduced by the preceding geometric modules.

9.1 Singular incidence towers

For persistent singular Pfaffian behaviour, the natural replacement for an orbit cap is an incidence tower. Starting from a rank stratum $Z_C^{(a)}$, one records successively

$$Z_C^{(a)} \leftarrow \mathcal{K}_C^{(a)} \leftarrow \mathcal{F}_C^{(a)} \leftarrow \mathcal{N}_C^{(a)} \leftarrow \mathcal{S}_C^{(a)}.$$

Here $\mathcal{K}$ records kernel spaces, $\mathcal{F}$ records kernel flags, $\mathcal{N}$ records conormal data, and $\mathcal{S}$ records syzygy or resolution data. If this tower becomes faithful, it reconstructs the local piece of C . If it does not, the residual symmetry is itself structural: field-extension, tensor-product, classical-polar, or primitive polar-core behaviour.

Remark 9.2. *The noetherian nature of the tower gives termination for any fixed prescribed refinement procedure, but it does not by itself prove that the tower is faithful in every primitive singular case. That faithfulness question is one of the precise residual problems isolated by the endpoint theory.*

10 Standard-structure reduction

The primitive polar-core material becomes mathematically clean when stated as a relative reduction. The residual group

$$\Gamma(C) = \text{Aut}(\text{Pol}(C)) \leq \text{PTL}(V)$$

acts on the projective geometry of V . If a small intrinsic base is found, landmark canonicalization applies. If a large residual symmetry remains, finite classical-group structure suggests standard geometric branches.

Definition 10.1 (Standard-structure primitives). *A standard-structure recognizer tests whether the residual polar-code automorphism group preserves a standard linear structure:*

subspace, direct sum, field extension, tensor product, subfield,

symplectic type, tensor induced, or classical form.

A standard-structure canonizer produces a canonical representative of C once such a structure has been identified. A bounded-base polar canonizer canonicalizes the remaining non-standard branch after individualizing a bounded-size base in the polar code.

Theorem 10.2 (Relative standard-structure reduction). *Assume available standard-structure recognizers and canonizers for the geometric branches above, and assume a bounded-base canonizer for the non-standard branch of the residual projective action. Then the isomorphism problem for special class-two exponent- p groups reduces to these primitives applied to*

$$C_\beta \leq \Lambda^2(G/Z(G))^*.$$

More precisely, the procedure

$$C \mapsto \text{Pol}(C) \mapsto \Gamma(C) \mapsto \text{standard branch or bounded-base branch} \mapsto \text{NF}_{\text{polar}}(C)$$

produces a complete normal form, provided each invoked recognizer and canonizer is correct.

Proof. The polar code determines C , so any complete canonicalization of the polar-code orbit is a complete canonicalization of the group. If a standard structure is recognized, the corresponding canonizer fixes the ambient gauge up to the normalizer of that structure and compares the induced components of C . The standard branches either reduce dimension, reduce to factors, reduce to a smaller field, or reduce the ambient group from $\mathrm{PGL}(V)$ to a classical or tensor-structured normalizer. If no geometric branch is present and the residual action has a bounded base, individualizing that base produces a projective frame after finitely many choices; landmark canonicalization then gives a normal form. Taking the lexicographic minimum over canonical choices removes residual finite ambiguity. Proposition 2.1 transfers the resulting normal form to the associated groups. $\square$

Remark 10.3. *Aschbacher's theorem and base-size results motivate the dichotomy: large residual subgroups of finite classical groups tend to preserve standard geometry, while non-standard primitive almost-simple actions have small bases in broad families. This note does not reprove those classification theorems and does not claim that the recognizers are automatically efficient. It isolates the algorithmic bottleneck: efficient recognition and canonical normalization of the standard polar-core branches.*

10.1 Standard branches

For reference, the branches have the following presentation-theoretic interpretations.

Subspace branch. A preserved subspace $0 < U < V$ decomposes

$$\Lambda^2 V^* \rightarrow \Lambda^2 U^*, \quad U^* \otimes (V/U)^*, \quad \Lambda^2 (V/U)^*.$$

The normal form is recursive on U , V/U , and the mixed term.

Direct-sum branch. For $V = V_1 \oplus \cdots \oplus V_t$,

$$\Lambda^2 V^* = \bigoplus_i \Lambda^2 V_i^* \oplus \bigoplus_{i < j} V_i^* \otimes V_j^*.$$

One canonicalizes the blocks and minimizes over the residual permutation action.

Field-extension branch. If V is a vector space over an extension K/k , then forms may be K -linear, semilinear, or obtained by trace. The field structure is visible through normalizer or adjoint data, and the problem is transferred to the smaller K -model with a Galois compatibility check.

Tensor-product branch. For $V = A \otimes B$,

$$\Lambda^2 (A \otimes B)^* \simeq (\Lambda^2 A^* \otimes \mathrm{Sym}^2 B^*) \oplus (\mathrm{Sym}^2 A^* \otimes \Lambda^2 B^*).$$

The polar rank and kernel incidences factor through the two tensor factors.

Subfield and tensor-induced branches. Subfield structure descends C to a minimal field of definition. Tensor-induced structure treats $V = V_0^{\otimes t}$ with the residual wreath-like permutation of factors.

Symplectic-type and classical-form branches. The residual group preserves an extraspecial or symplectic-type structure, or a nondegenerate orthogonal, unitary, or symplectic form. The polar core is then canonicalized as a substructure of the corresponding polar space.

Non-standard branch. When no standard geometry is present and a bounded base is available, polar-base individualization gives a finite canonical normal form with cost controlled by the base size.

11 Resolution spectrum for the class-two papers

The surrounding group-theoretic papers fit into the following spectrum.

Quadratic word calculus. In class two and exponent p , words reduce to linear exponent sums and alternating commutator tensors. This gives a normal form for word observables and explains why arity is a genuine observable budget.

Scalar word measures. Ordinary word-map distributions factor through R_C . Pfaffian representations of the same smooth plane curve can have the same R_C but yield non-isomorphic groups.

Support-image and star profiles. Star probabilities, rank-support enumerators, and support-image profiles record contraction-rank geometry. They refine scalar rank data but are still localized observables.

Exact pencil thresholds. For regular alternating pencils, depth is governed by spectral divisors. An irreducible block of degree m has CSI depth m and word depth $2m$.

Fourier-tagged pencil tomography. For $\dim W = 2$, explicit Fourier-tagged commutator-word distributions recover the full skew-symmetric Kronecker data. Thus the polar-code endpoint is reached by a finite family of group-theoretic observables in derived rank two.

Odd-Pfaffian reconstruction. In higher derived rank and odd $\dim V$, pointwise Fourier-tagged moments can recover Pfaffian degeneracy schemes, and on a rigid open locus those schemes recover the tensor.

Kernel incidence. For Pfaffian word-measure twins, kernel incidence sees the vector-bundle data invisible to scalar word measures. It is a concrete intermediate layer between scalar rank and the full polar code.

Polar code. The polar code is complete for all alternating systems. It gives the endpoint normal form and clarifies the residual hard case: efficient canonicalization of primitive polar cores.

12 Conclusion

Presentation theory turns the group-isomorphism problem here into a fibre-resolution problem. The weak observable

ordinary word maps

collapses to scalar rank. Intermediate observables resolve specific fibres with exact costs: support-image profiles for pencil thresholds, Fourier-tagged moments for Kronecker data, Pfaffian schemes for generic odd systems, and kernel incidence for Pfaffian moduli. The full polar code resolves the entire fibre:

$$\text{Pol}(C_\beta) \iff C_\beta \iff G_\beta.$$

Thus the conceptual progression is

$$\begin{aligned} \text{rank profile} &\rightarrow \text{localized support data} \rightarrow \text{tomographic word moments} \\ &\rightarrow \text{kernel incidence} \rightarrow \text{polar code.} \end{aligned}$$

The new structural point is that the endpoint is elementary and complete, while the hard mathematical work lies in measuring the cost of reaching it from weaker, more natural, or more computable observables.

The complete invariant is not the final algorithmic story. The refined endpoint is:

word measures collapse to scalar rank; polar code completes the invariant;

generic systems are reconstructed geometrically;

worst-case systems reduce to standard geometry or bounded bases.

The precise remaining problem is to make the standard-structure recognizers and canonizers efficient enough to compete with the best matrix-space isometry and class-two p -group isomorphism algorithms.

References

- [1] M. Aschbacher, On the maximal subgroups of the finite classical groups, *Invent. Math.* 76 (1984), 469–514.
- [2] R. Baer, Groups with abelian central quotient group, *Trans. Amer. Math. Soc.* 44 (1938), 357–386.
- [3] P. A. Brooksbank, Y. Li, Y. Qiao, and J. B. Wilson, Improved algorithms for alternating matrix space isometry, *LIPICs ESA* 173 (2020), 26:1–26:15.
- [4] D. Buchsbaum and D. Eisenbud, Algebra structures for finite free resolutions, and some structure theorems for ideals of codimension 3, *Amer. J. Math.* 99 (1977), 447–485.
- [5] T. C. Burness, On base sizes for almost simple primitive groups, arXiv:1803.10955.
- [6] W.-L. Chow, On the geometry of algebraic homogeneous spaces, *Ann. of Math.* 50 (1949), 32–67.
- [7] S. R. Ghorpade and K. V. Kaipa, Automorphism groups of Grassmann codes, *Finite Fields Appl.* 23 (2013), 80–102.
- [8] R. M. Guralnick and R. Lawther, Generic stabilizers in actions of simple algebraic groups II: higher Grassmannian varieties, arXiv:1904.13382.
- [9] X. Sun, Faster isomorphism for p -groups of class 2 and exponent p , arXiv:2303.15412.
- [10] J. B. Wilson, Decomposing p -groups via Jordan algebras, *J. Algebra* 322 (2009), 2642–2679.
- [11] C. Zhang, Faster isomorphism testing of p -groups of Frattini class 2, FOCS 2024.