A FAST ISOMORPHISM TEST FOR GROUPS WHOSE LIE ALGEBRA HAS GENUS 2

PETER A. BROOKSBANK, JOSHUA MAGLIONE, AND JAMES B. WILSON

Abstract. Motivated by the desire for better isomorphism tests for finite groups, we present a polynomial-time algorithm for deciding isomorphism within a class of p-groups that is well-suited to studying local properties of general groups. We also report on the performance of an implementation of the algorithm in the computer algebra system magma.

Contents

2010 *Mathematics Subject Classification.* 20D15, 20D45, 15A22, 20B40.
Key words and phrases. group isomorphism, pairs of forms, Pfaffian, adjoint tensor.

1. Introduction

The best known general algorithms to test whether a pair of finite groups of order n are isomorphic use $n^{O(\mu(n))}$ operations, where $\mu(n) = \max\{e_i : n = p_1^{e_1} \cdots p_s^{e_s}, p_i \text{ prime } \}$. These algorithms are far too slow for most practical purposes, and their complexity (super-polynomial in the order of the groups) falls short of natural theoretical benchmarks. Significant theoretical or practical progress seems beyond the reach of current methods.

A new strategy developed by the first and third authors in collaboration with E.A. O'Brien breaks up the work into multiple overlapping instances of isomorphism of p-groups with small commutator subgroups. This raises the question of how non-abelian a p-group can be and still have a highly efficient isomorphism test. The results of this paper show that, for finite p-groups whose commutator subgroup is central and isomorphic to $\mathbb{Z}_p \times \mathbb{Z}_p$, there is an $O(p^3 + (\log n)^{2\omega})$ time test for isomorphism, where $2 \leqslant \omega < 3$ is the exponent of efficient matrix multiplication [vzGG, Chapter 12]. Note, it takes at least $O((\log n)^3)$ bits of information to even describe a p-group of order n.

To state our main theorem, we recall a notion from Lie theory. In [K], Knebelman defines the *genus* of a Lie K-algebra to be $\dim_K L - d(L)$, where $d(L)$ is the size of the smallest generating set for L. Every finite p-group, G, has an associated Lie $\mathbb{Z}_p$-algebra, $L(G)$, and there is a unique largest extension $K : \mathbb{Z}_p$ such that $L(G)$ is a Lie K-algebra. So we define the *genus* of G to be the genus of $L(G)$ as a K-algebra.

We assume that groups are input succinctly, for example using permutations, matrices, or polycyclic presentations; see Section 2.3 for a more detailed discussion. Treating the cost of operations in such groups as constant, we prove the following.

Theorem 1.1. *There are algorithms that, given groups G_1 and G_2 of order p^m with derived subgroups of order p^s,*

(a) decide if the G_i are p-groups of class 2, exponent p, and genus 2, and if so
(b) return the coset of isomorphisms $G_1 \to G_2$ (or the empty set if none exist).

The algorithm for part (a) runs in time $O(m^{2\omega} \log^2 p)$. For part (b) we let t bound the number of pairwise non-isomorphic central product factors of G_1 or G_2 and prove the algorithm runs in time $O((m^{2\omega-2} + \min\{sp^{3s/2}, t!\})m^2 \log^2 p)$.

1.1. Implementation.

We have implemented the algorithms of Theorem 1.1 in the computer algebra system MAGMA [BCP]. The implementation is available upon request from the authors. A detailed analysis of its performance is given in Section 8, but we summarize here the results of one experiment to illustrate its efficiency. We constructed 1557 random 5-groups of genus 2 having orders ranging from 5^5 to 5^{256}, and generated for each a random isomorphic copy. We then tasked our implementation to find an explicit isomorphism between each pair of groups, plotting the completion time on a graph. Figure 1.1 shows the performance. We intended to compare the performance of our implementation with that of existing functions in MAGMA, but even for groups of order 5^7 the existing routines exhausted the 500GB of memory on our largest machine. Instead Figure 8.1 shows how our algorithm compares with the cost of solving random systems of linear equations in approximately d^2 variables and equations, where $d = \log_5 |G|$.

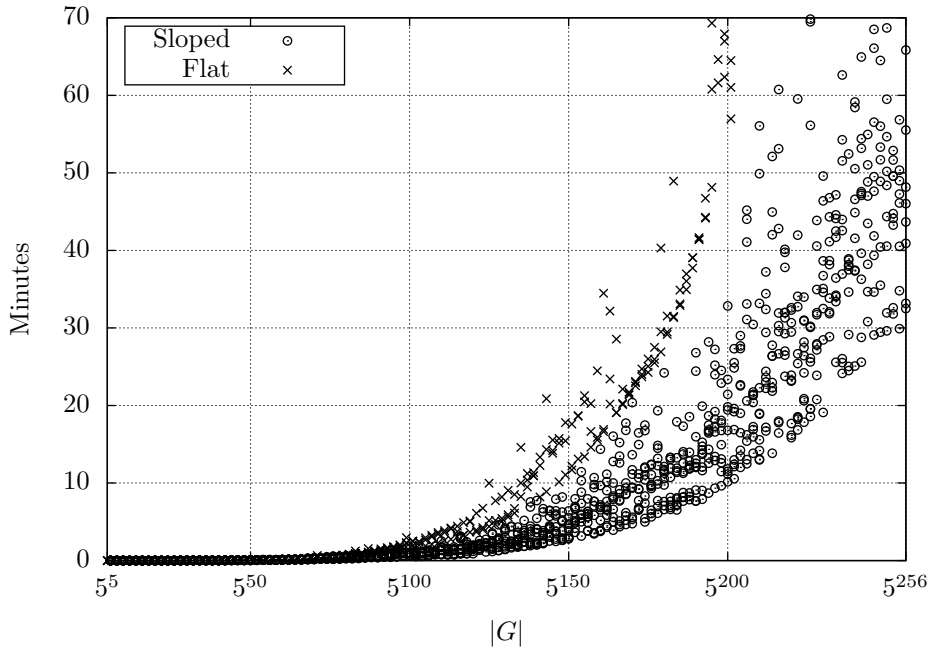

FIGURE 1.1. Performance data for tests to confirm isomorphism between groups of order 5^{d+2} for increasing d. Sloped and flat designate further properties of the input discussed in Theorem 1.2. See Figure 8.1 for a comparison of how closely our algorithm tracks with the speed of MAGMA's methods for solving systems of linear equations of comparable size.

1.2. **Classification problems.** Groups of genus 2 were first studied in the context of the *finite-tame-wild* trichotomy theorems. In particular, they have been shown to lie on the tame-wild boundary. General p-groups of genus 2 are *wild* (their classification would imply a classification of all finite-dimensional algebras) [BLS, BDL$^+$]. Those of exponent p are, however, *tame* [V1] which means they decompose, via central products, into one-parameter families. A classification of these one-parameter families for arbitrary fields was not previously known; to prove Theorem 1.1 we required such a classification.

Theorem 1.2. *A centrally indecomposable p-group of exponent p, and of genus 2 over a field k, is isomorphic to one of the following two types of groups:*

(i) a quotient by a central subgroup N of a Heisenberg group,

$$H = \left\{ \begin{bmatrix} 1 & e & z \\ 0 & 1 & f \\ 0 & 0 & 1 \end{bmatrix} : \begin{array}{c} e, f, z \in k[x]/(a(x)^c), \\ a(x) \text{ irreducible} \end{array} \right\},$$

where as matrices $1 - N$ is a k-subspace of codimension 2 in $1 - [H, H]$; or

(ii) the matrix group

$$H^\flat = \left\{ \left[\begin{array}{c|ccccc|c} I_2 & e_1 & \cdots & e_m & 0 & z_1 \\ & 0 & e_1 & \cdots & e_m & z_2 \\ \hline & & & & & f_0 \\ & & & I_{m+1} & & \vdots \\ & & & & & f_m \\ \hline & & & & & 1 \end{array}\right] : e_i, f_j, z_\ell \in k \right\}.$$

1.3. **Overview.** The reader familiar with p-groups may already have noticed the connection, which will be further elucidated in Sections 2 and 3, between groups of genus 2 and pairs $\{\Phi_1, \Phi_2\}$ of alternating forms over a finite field $\mathbb{F}_q$. Indeed, we use the classification of such pairs by Bond [B3] and Scharlau [S1] (who use an earlier classification of pairs of matrices – or *Kronecker modules* as they are known – by Kronecker and Dieudonné [D]). We also exploit a recently discovered Galois connection between adjoints of bilinear maps and tensor products [W2, BW2] to prove Theorem 1.2.

By itself, Theorem 1.2 is not sufficiently powerful to decide isomorphism among genus 2 groups (not even if we restrict to centrally indecomposable groups). There exist non-isomorphic groups of genus 2 whose centrally indecomposable factors are isomorphic (see Example 3.16). Hence, theorems of Krull-Remak-Schmidt type, upon which the classification of Kronecker modules depends, simply do not exist for groups of genus 2. Nevertheless, we prove a transitivity result on fully-refined central decompositions of groups of genus 2 (Theorem 3.15). This makes for a well-defined generalization of the Pfaffian to pairs of alternating forms. The resulting characterization of isomorphism classes, presented in Theorem 3.22 in terms of bilinear maps, leads to an isomorphism test that is effective when $\mathbb{F}_q$ is small.

When $\mathbb{F}_q$ is large, we use a general technique for isomorphism testing in groups of p-class 2. Dubbed the *adjoint-tensor method*, this technique was proposed in [BW2] by the first and third authors as a means of bridging the gap between the generic but typically slow general algorithms, and incredibly fast but highly specialized isomorphism tests such as the one in [LW]. The adjoint-tensor method is presented in a mildly restricted form in Section 4. It requires the user to solve several problems – such as *algebra conjugacy*, *algebra normalizer*, and *subspace transporter* – that are known in their general forms to be hard. The main work is to show that in our particular setting each of these three problems has an efficient solution. The details of the test, which include effective methods for computing with certain quotients of the notorious Nottingham group, are presented in Sections 5 through 7.

We discuss our MAGMA implementation in Section 8, and give further details on its performance. Section 9 contains some concluding remarks, including observations about the number of groups of genus 2.

2. Nilpotent Groups and Bimaps

We describe the relationship between groups of nilpotence class 2 and bilinear maps. This has a long history going back to work of Brahana and Baer in the 1930's. Henceforth, all groups are finite.

2.1. **Bimaps.** Let k be a commutative ring, and U, V, W (left) k-modules. A *k-bilinear map*, which we abbreviate to *k-bimap*, is a function $\circ\colon U \times V \rightarrowtail W$ such that, for all $u, u' \in U$, $v, v' \in V$, and $\alpha \in k$,

$$(u + \alpha u') \circ v = u \circ v + \alpha(u' \circ v),$$
$$u \circ (v + \alpha v') = u \circ v + \alpha(u \circ v').$$

The *radicals* of $\circ$ are $U^{\perp} = \{v \in V\colon U \circ v = 0\}$, $V^{\top} = \{u \in U : u \circ V = 0\}$, and $W^{+} = W/(U \circ V)$. We say $\circ$ is *fully-nondegenerate* if all three radicals are trivial. If $U = V$ and $v \circ v = 0$ for all v, then $\circ$ is *alternating*. We reserve the use of U, V and W for these three variables of a bimap and write $U_{\circ}$, $U_{\bullet}$, and so forth if we need to distinguish between these components for separate bimaps $\circ$, $\bullet$.

A *homotopism* between bimaps $\bullet\colon U_{\bullet} \times V_{\bullet} \rightarrowtail W_{\bullet}$ and $\circ\colon U_{\circ} \times V_{\circ} \rightarrowtail W_{\circ}$ is a triple $f = (f_U, f_V; f_W) \in \mathrm{Hom}(U_{\bullet}, U_{\circ}) \times \mathrm{Hom}(V_{\bullet}, V_{\circ}) \times \mathrm{Hom}(W_{\bullet}, W_{\circ})$ such that

$$(\forall u \in U_{\bullet}, \forall v \in V_{\bullet}) \qquad\qquad uf \circ vf = (u \bullet v)f.$$

When working with such a triple of maps, writing uf for $u \in U$ means uf_U, whereas vf for $v \in V$ means vf_V, and so on. Bimaps with homotopisms form a natural category [W2]. A homotopism in which all maps are invertible is an *isotopism*. We typically work here with alternating bimaps, and for such bimaps we shall further insist that $f_U = f_V$ and refer to an isotopism between $\bullet$ and $\circ$ as a *pseudo-isometry*.

When we need to describe a bimap in an example – or for computation – we do so via matrices. Fix generating sets X, Y, Z for U, V, W, respectively, as k-modules. For $x \in X$, $y \in Y$, there exist $\alpha_{xyz} \in k$ ($z \in Z$) such that

$$x \circ y = \sum_{z \in Z} \alpha_{xyz} z.$$

The scalars α_{xyz} are called *structure constants* of $\circ$ relative to X, Y, Z. When k is a field, these constants are uniquely determined by the choices of X, Y, Z and we record the data using matrices $\Phi_z = [[\alpha_{xyz}]]$, where $z \in Z$ and each Φ_z is an $|X| \times |Y|$ matrix. When $\circ$ is alternating, each Φ_z represents an alternating form on $U = V$, and $\{\Phi_z\colon z \in Z\}$ is commonly known as a *system of forms* [BF, BO].

2.2. **Isoclinism and isomorphism of groups.** One can associate to each nilpotent group G of class 2 an alternating bimap $\circ_G$. Equivalence of such bimaps up to pseudo-isometry corresponds to an equivalence of groups that is in general weaker than isomorphism. This equivalence was introduced by Philip Hall [H] and is known as *isoclinism*. The relationship between isoclinism and isomorphism for groups is akin to that between homotopy equivalence and homeomorphism for topological spaces.

Commutation in a group is a function $[,]\colon G \times G \to G$ whose image is not, in general, a subgroup of G. However, the subgroup generated by this image is the *commutator subgroup* and is denoted $[G, G]$ or G'. Commutation is also not a homomorphism and hence has no kernel. However, the *center* of G, namely $Z(G) = \{g \in G\colon [G, g] = [g, G] = 1\}$ consists of those elements that do not influence the outcome of commutation. Thus, given M such that $G' \leqslant M \leqslant Z(G)$ we can reduce $[,]$ to a *commutation word map*

$$\begin{array}{rccc} \bullet_{G,M}\colon & G/M \times G/M & \to & G' \\ & (\bar{x}, \bar{y}) & \mapsto & [x, y], \end{array}$$

where $\bar{x}$ denotes the coset xM. If $M = Z(G)$ we write $\circ_G$. Comparing groups G and H only up to their commutation structures is therefore comparing the maps $\circ_G$ and $\circ_H$. Doing so requires homomorphisms $f\colon G/Z(G) \to H/Z(H)$ and $\hat{f}\colon G' \to H'$ such that

$$(\forall x, y \in G) \qquad \bar{x}f \circ_H \bar{y}f = (\bar{x} \circ_G \bar{y})\hat{f}.$$

The pair $(f, \hat{f})$ is a *homoclinism* and, if the pair is invertible, it is an *isoclinism*.

In [B1], Baer established a fundamental correspondence for class 2 nilpotent groups that may already be evident from the foregoing discussion.

Theorem 2.1 (Baer correspondence)**.** *If $[G, G] \leqslant Z(G)$ then $\circ_G$ is a fully-nondegenerate alternating $\mathbb{Z}$-bimap. Also, two groups G and H of nilpotence class 2 are isoclinic if, and only if, $\circ_G$ and $\circ_H$ are pseudo-isometric.*

The next crucial observation follows from the Universal Coefficients Theorem applied to group cohomology. (Direct proofs are also known; see, for example, [W1, Proposition 3.10].)

Proposition 2.2. *Two p-groups of nilpotence class 2 and exponent p and the same order are isoclinic if, and only if, they are isomorphic.*

Theorem 2.1 and Proposition 2.2 lead us to study the *pseudo-isometry* group of an alternating bimap $\circ\colon V \times V \rightarrowtail W$, namely

$$\Psi\mathrm{Isom}(\circ) = \{(\varphi, \hat{\varphi}) \in \mathrm{Aut}(V) \times \mathrm{Aut}(W)\colon \forall u, v \in V, u\varphi \circ v\varphi = (u \circ v)\hat{\varphi}\}.$$

If $\circ$ is k-bilinear over a field k, we can further separate semilinear and linear pseudo-isometries $\Psi\mathrm{Isom}_k(\circ) = \Psi\mathrm{Isom}(\circ) \cap \mathrm{GL}_k(V) \times \mathrm{GL}_k(W)$ with exact sequence

$$1 \longrightarrow \Psi\mathrm{Isom}_k(\circ) \longrightarrow \Psi\mathrm{Isom}(\circ) \longrightarrow \mathrm{Gal}(k).$$

The following observation allows us to focus on k-linear isometries.

Lemma 2.3. *Let $\circ, \bullet\colon V \times V \rightarrowtail W$ be fully-nondegenerate and k-bilinear, and $(\varphi, \hat{\varphi})$ a k-semilinear pseudo-isometry of $\circ$. If $\sigma, \hat{\sigma} \in \mathrm{Gal}(k)$ are such that*

$$(\forall s \in k, \forall u \in V, \forall w \in W) \qquad (su)\varphi = s^{\sigma}(u\varphi) \text{ and } (sw)\hat{\varphi} = s^{\hat{\sigma}}(w\hat{\varphi}),$$

then $\sigma = \hat{\sigma}$.

Proof. We can test this on the generators $u \circ v$ of W. As $\hat{\varphi}$ is an isomorphism of W, $\{(u \circ v)\hat{\varphi} : u, v \in V\}$ generates W. So $s^{\hat{\sigma}}((u \circ v)\hat{\varphi}) = (su)\varphi \circ v\varphi = s^{\sigma}((u \circ v)\hat{\varphi})$. □

For a group G, we denote by $\mathrm{Aut}(G)$ its group of automorphisms, and by $C_{\mathrm{Aut}(G)}(G/M)$, for any $M \lhd G$, the subgroup of automorphisms that induce the identity on G/M. Questions of isomorphisms and automorphisms of groups reduce to ones about bimaps as follows (cf. [W1, Proposition 3.8]).

Proposition 2.4. *Let G be a p-group of class 2, $V = G/G'$ and $W = G'$, so that $\bullet_G\colon V \times V \rightarrowtail W$. Then the following are exact sequences*

$$1 \longrightarrow C_{\mathrm{Aut}(G)}(V) \xrightarrow{\iota} \mathrm{Aut}(G) \xrightarrow{\pi} \Psi\mathrm{Isom}(\bullet_G),$$

and, with $R = Z(G)/G'$,

$$1 \to \mathrm{Hom}_{\mathbb{Z}_p}(V, R) \to \Psi\mathrm{Isom}(\bullet_G) \to \Psi\mathrm{Isom}(\circ_G) \times \mathrm{Aut}(R) \to 1.$$

If G has exponent p then π is surjective and split, and $C_{\mathrm{Aut}(G)}(V) \cong \mathrm{Hom}(V, W)$.

Finally, for a fixed alternating bimap $\circ\colon V \times V \rightarrowtail W$, the *isometry* group is

$$\mathrm{Isom}(\circ) = \{\varphi \in \mathrm{Aut}(V)\colon \forall u, v \in V, u\varphi \circ v\varphi = u \circ v\}.$$

This is the kernel of the restriction of $\Psi\mathrm{Isom}(\circ)$ to W. Note that if $\circ$ is nondegenerate then $\mathrm{Isom}_k(\circ) = \mathrm{Isom}(\circ)$ because $(su)\varphi \circ v\varphi = s(u \circ v) = (s(u\varphi)) \circ v\varphi$.

2.3. **Computational models for groups.** Efficient algorithms exist to determine crucial information about groups. Details and proofs can be found in [S3, HEO]. The meaning of efficiency depends on how groups are specified for computation.

The pioneering work of Sims, Cannon, and Neubüser in the 1960s and 1970s led to the standard models of computation that we use today. It is most common to specify general finite groups by small sets of generators (matrices over finite fields or permutations of a finite set). Special classes of groups admit certain types of structured presentations as feasible computational models. For example, polycyclic presentations are often used for computations with solvable groups. Algorithms for p-groups should certainly apply to these specialized models but we caution that the complexity of multiplication can be exponential [LGS1, LGS2].

The notion of a "black-box" group was introduced by Babai and Szemerédi [BS] in order to strip away information specific to the particular representation, and thereby force algorithms to deal only with the algebraic structure of the group. All of our algorithms apply to groups G where algorithms exist to solve the following tasks:

(i) find $|G|$;
(ii) given $x \in G$ and a sequence $x_1, \ldots, x_c \in G$, either write x as a word over $x_1, \ldots, x_c \in G$, or else prove $x \notin \langle x_1, \ldots, x_c\rangle$;
(iii) find generators for $Z(G)$ and for G';
(iv) decide if G is nilpotent of class 2; and
(v) if G is nilpotent of class 2, construct a system of forms for $\circ_G$.

While it is standard to represent groups as permutations of a set, we note that often p-groups cannot be represented faithfully as permutations on small sets. As an illustration, P. Neumann [N] showed that extraspecial groups $2{+}_s^{2m+1}$ have no faithful permutation representation on fewer than 2^m points. However, Neumann's groups are quotients of D_8^m, and so they can be represented as quotients of permutation groups on $4m$ points. (This phenomenon occurs also for p-groups of odd prime power – see, for example, Proposition 9.2.) To address this issue, one can choose to work instead with the *permutation group quotient* model proposed by Kantor and Luks [KL]. We note that all of the necessary foundational computations can be carried out effectively in this model:

Proposition 2.5. *Given a group G as a quotient of a permutation group, in polynomial time one can solve all of the problems listed in (i) through (v) above.*

Proof. For (i)–(iv) see [S3, Chapter 6] and [KL]. For (v), fix bases $\{x_1, \ldots, x_d\}$ and $\{w_1, \ldots, w_e\}$ for the abelian groups $G/Z(G)$ and G', respectively. The structure constants for the associated system of forms are obtained by writing each $[x_i, x_j]$ as a vector relative to $\{w_1, \ldots, w_e\}$. □

We shall state and prove various results for bimaps that require us to work with large fields. We therefore allow ourselves to factor polynomials using randomized

Las Vegas polynomial-time factorization algorithms. (A *Las Vegas algorithm* always returns a correct result but with small, user prescribed, probability reports failure.) Such methods can always be "derandomized" whenever the characteristic p is bounded by the input size – as is the case with permutation group quotients. We refer the reader to [vzGG] for further information on these matters.

Complexity of group isomorphism. Group isomorphism is often reported as having complexity $n^{\log n + O(1)}$, where n is the order of the input groups, in part because it was reported this way in an influential paper by Miller [M1]. A more accurate description of the simple bound is that it takes time $n^{d+O(1)}$, where d is the common size of a smallest generating set for the input groups. Guralnick and Lucchini have independently shown that d is bounded by $\mu(n) + 1$ [BNV, p.146]. Rosenbaum and Wagner [RW] show that the leading constant in the exponent can be decreased below 1 in many circumstances. There are also numerous unanalyzed improvements in the literature [ELGO, O] that likely influence the cost. The estimate of $n^{O(\mu(n))}$ for some leading constant less than 1 is a reasonable over estimate of the best bound by today's methods.

3. Groups of Genus 2

In this section we propose an integral metric for the "complexity" of a nilpotent group. Inspired by an analogous metric introduced by Knebelman [K] to measure the complexity of Lie algebras, we call this number the *genus* of a group.

3.1. **The centroid and genus of a group.** In Section 1 we defined the genus of a p-group in terms of its associated Lie algebra using Knebelman's original definition. Here, we give an equivalent formulation using bimaps that is better suited to our computational goals. Let k be a commutative ring, and $\circ\colon U \times V \rightarrowtail W$ a k-bimap. The *centroid* of $\circ$ is the largest ring, C, over which $\circ$ is C-bilinear, namely

$$C(\circ) = \{\sigma \in \mathrm{End}(U) \times \mathrm{End}(V) \times \mathrm{End}(W)\colon \forall u, \forall v, (u\sigma) \circ v = (u \circ v)\sigma = u \circ (v\sigma)\}.$$

This explicit definition of the ring makes it clear that $C(\circ)$ may be obtained as the solution of a system of linear equations. It is understood that $\sigma \in C(\circ)$ acts naturally on U, V, and W but we can write $\sigma = (\sigma_U, \sigma_V; \sigma_W)$ if we wish to clarify the action on the individual k-modules. If $\circ$ is fully-nondegenerate – as is the case with the commutation bimap of a group – then $C(\circ)$ is commutative. The following connection between centroids and direct products was proved in [W4, Section 6].

Theorem 3.1. *A finite nilpotent group G of class 2 is isoclinic to a direct product of proper nontrivial subgroups if, and only if, the centroid $C(\circ_G)$ is a direct product of proper subrings.*

Being concerned with questions of isomorphism, we focus on directly indecomposable groups. If G is such a group, by Theorem 3.1, $C = C(\circ_G)$ is a local ring. Thus, if $J = J(C)$ is the Jacobson radical of C, W/WJ is a vector space over the residue field C/J, and we define the *rank* of W to be the dimension of this space.

Definition 3.2. Let G be a nilpotent group of class 2. Then G is isoclinic to a direct product $H_1 \times \cdots \times H_s$ of directly indecomposable groups. The *genus* of G is maximum rank of $[H_i, H_i]$ as a $C(\circ_{H_i})$-module.

The concept of genus arose first in Knebelman's attempts to classify Lie algebras and general nonassociative algebras [K]. He observed that when the dimension of a Lie k-algebra L was close to the minimum number, $d(L)$, of generators, there are relatively few variable relations. Accordingly, he proposed that algebras of low *genus* – which he defined as $\dim L - d(L)$ – should be easier to classify. For instance, if L is abelian then $\dim L - d(L) = 0$, and if L is a Heisenberg Lie algebra then $\dim L - d(L) = 1$. Later, Bond tackled the classification of Lie algebras of genus 2, and reduced the problem to the class 2 nilpotent Lie algebras of genus 2 [B3]. The latter problem remains very difficult. In fact the classification of 6-dimensional Lie algebras has only recently been completed [M2, CdGS], and the nilpotent genus 2 cases are the most involved.

3.2. **Some groups of low genus.** To reveal some important subtleties in the definition of genus, and to provide concrete examples of the groups we propose to study, we introduce some groups of genus 1 and genus 2.

(a) Every group with cyclic central commutator subgroup has genus 1. For such G with $\circ_G\colon G/Z(G) \times G/Z(G) \rightarrowtail \mathbb{Z}_m$ we have $C(\circ_G) = \mathbb{Z}_m = \mathbb{Z}_{p_1^{e_1}} \oplus \cdots \oplus \mathbb{Z}_{p_s^{e_s}}$, with each p_i a distinct prime. As G is a direct product of its Sylow subgroups, we need only the maximum genus when restricted to each p_i. As each $\mathbb{Z}_{p_i^{e_i}}$ is cyclic, the genus of each Sylow subgroup is 1.

(b) Any group with central commutator subgroup isomorphic to $\mathbb{Z}_m \times \mathbb{Z}_n$ has genus at most 2. Let G be such a group. If $(m, n) = 1$, then $\mathbb{Z}_m \times \mathbb{Z}_n$ is cyclic and G has genus 1. Else, G is a product of Sylow subgroups. Let P be a Sylow p-subgroup of G of largest genus. We may assume $P' \cong \mathbb{Z}_{p^e} \times \mathbb{Z}_{p^f}$, $e \geqslant f \geqslant 1$. Either $C(\circ_P) \cong \mathbb{Z}_{p^e}$ (in which case P is genus 2), or $C(\circ_P)$ is not local and P is isoclinic to a nontrivial direct product (so P is genus 1).

(c) Fix any commutative Artinian ring K, and consider the Heisenberg groups

$$H_m(K) = \left\{ \begin{bmatrix} 1 & u & s \\ 0 & I_m & v^{\mathrm{tr}} \\ 0 & 0 & 1 \end{bmatrix} : u, v \in K^m, s \in K \right\}.$$

If $K = K_1 \oplus \cdots \oplus K_s$ is a decomposition of K into local rings, then

$$H_m(K) \cong H_m(K_1) \times \cdots \times H_m(K_s),$$

so the genus is the maximum genus of any $H_m(K_i)$. The bimap of $H_m(K_i)$ is simply the alternating nondegenerate form $K_i^{2m} \times K_i^{2m} \rightarrowtail K_i$ having K_i as centroid. Since K_i is commutative and local, $K_i/J(K_i)$ is a field, so $H_m(K_i)$ has genus 1. Hence, all Heisenberg groups have genus 1.

While all of these examples are somewhat elementary, from a classification perspective we have already entered turbulent waters. For instance, classifying the groups with cyclic central commutator subgroup in part (a) has taken the combined work of several authors including Leong [L1], Finogenov [F], and Blackburn [B2]. The Heisenberg groups in (c) were only recently characterized in abstract terms (with no a priori knowledge of m or K) for the case when K is a field [LW, Theorem 3.1]. (Our results extend that characterization to K an arbitrary cyclic algebra.)

It may surprise the reader that groups seeming to have genus $g > 1$ are in fact genus 1. For example, if $[K\colon \mathbb{Z}_p] = g$, then $H_1(K)$ is a group whose central commutator subgroup is isomorphic to $\mathbb{Z}_p^g$, so it would seem that one can easily build a group of high genus. However, the centroid recovers field structure and, viewed

as a vector space over the centroid, the commutator subgroup is 1-dimensional. A direct product $H_1(\mathbb{Z}_p)^g$ also has commutator subgroup $\mathbb{Z}_p^g$. Via Theorem 3.1 and Definition 3.2, however, those examples are again genus 1. Note, moreover, that our definition of centroid is invariant under extensions: if G is a group of genus g over a field k and H is a group such that $\circ_H$ is the tensor of $\circ_G$ with a field extension K of k, then H has genus g over K.

3.3. **Central decompositions, hyperbolic pairs, and adjoints.** The groups of genus 2 admit two important decompositions. The first decomposes the group as a central product of subgroups, and the second as a product of two abelian normal subgroups whose intersection is central. We shall make essential use of both types of decomposition in our algorithms, so we now introduce them and give characterizations that facilitate effective computation.

Definition 3.3. A *central decomposition* of a group G is a set, $\mathcal{H}$, of subgroups generating G such that for $H \in \mathcal{H}$, $G \neq \langle \mathcal{H} - \{H\}\rangle$ and $[H, \langle \mathcal{H} - \{H\}\rangle] = 1$. We say that G is *centrally indecomposable* if $\{G\}$ is its only central decomposition.

A detailed treatment of central decompositions of p-groups is the subject of [W1], and we shall use some of the results therein. The second decomposition we need mimics hyperbolic pairs in the sense of symplectic geometry. It was introduced in [LW, Section 6] to work with 2-groups, but we use it here for arbitrary p-groups.

Definition 3.4. A *hyperbolic pair* for a group G is a pair M, N of normal abelian subgroups of G such that $G = MN$ and $M \cap N \leqslant Z(G)$.

Both central and hyperbolic decompositions may be obtained from a ring that is easily computed from $\circ_G$, namely the ring of adjoints. In a similar vein to our definition of centroid, we introduce the *adjoint ring*, $A(\circ)$, of a bimap $\circ\colon U \times V \rightarrowtail W$ as the largest ring, A, over which $\circ$ factors through $U \otimes_A V$, namely

$$A(\circ) = \{\mu \in \mathrm{End}(U) \times \mathrm{End}(V)^{\mathrm{op}} : \forall u \forall v, u\mu \circ v = u \circ \mu v\}.$$

Again, $A(\circ)$ may be obtained as the solution of a system of linear equations [W5, BW4, BW3]. As $\mathrm{End}(V)^{\mathrm{op}}$ suggests, we find it convenient to work with the opposite ring in the second component – thus A acts on U on the right but on V on the left. If we need to clarify the action we write $u\mu = uL_\mu$ and $\mu v = vR_\mu$.

If $\circ\colon V \times V \rightarrowtail W$ is a nondegenerate, alternating bimap, then $A(\circ)$ is faithfully represented in $\mathrm{End}(V)$ and in $\mathrm{End}(V)^{\mathrm{op}}$. This endows $A(\circ)$ with a natural anti-isomorphism interchanging L_μ and R_μ, giving it the structure of a $*$-ring. The connections to central decompositions and hyperbolic pairs come from the existence of certain types of idempotents in this $*$-ring. We say that an idempotent, e, in $A(\circ)$ is *self-adjoint* if $e^* = e$, and *hyperbolic* if $e^* = 1 - e$. Recall that idempotents e, f in a ring are *orthogonal* if $ef = 0 = fe$.

Lemma 3.5. *A finite nilpotent group, G, of nilpotence class 2 has*

(i) a central decomposition $\{H_1, \ldots, H_s\}$ if, and only if, $A(\circ_G)$ has a set $\{e_1, \ldots, e_s\}$ of pairwise orthogonal, self-adjoint idempotents that sum to 1, and
(ii) a hyperbolic pair if, and only if, $A(\circ_G)$ has hyperbolic idempotents.

Proof. A proof of (i) may be found in [W1, Theorem 4.10].

For (ii), let $Z = Z(G)$, and $\circ = \circ_G\colon G/Z \times G/Z \rightarrowtail G'$. Suppose M, N is a hyperbolic pair for G, and put $E = MZ/Z$ and $F = NZ/Z$. Since $G = MN$ and

$M \cap N \leqslant Z(G)$ we have $V = G/Z = E \oplus F$. Let e denote the projection idempotent onto E with kernel F. Hence, $1-e$ is the projection idempotent onto F with kernel E. As M and N are abelian, note $E \circ E = 0 = F \circ F$, so for all $u, v \in V$,

$$\begin{aligned} ue \circ v &= ue \circ ev + ue \circ (1-e)v \\ &= u \circ (1-e)v - u(1-e) \circ (1-e)v = u \circ (1-e)v. \end{aligned}$$

In particular, $e \in A(\circ)$, and $e^* = 1 - e$.

Conversely, observe that if $e \in A(\circ)$ and $e^* = 1 - e$, then $V = Ve \oplus V(1-e)$ and $Ve \circ eV = Ve(1-e) \circ V = 0$. Hence, $M = \{g \in G\colon (gZ)e = gZ\}$ and $N = \{g \in G\colon gZ(1-e) = gZ\}$ is a hyperbolic pair for G. □

We say that a central decomposition is *fully-refined* if each term in the decomposition is centrally indecomposable. Lemma 3.5 is the tool we need to compute such decompositions, and also hyperbolic pairs.

Theorem 3.6. *There are polynomial-time algorithms for each of the following:*

(a) construct a fully-refined central decomposition of a given finite p-group; and
(b) decide if a given p-group has a hyperbolic pair and construct one if it does.

Proof. A proof of (a) may be found in [W5, Theorem 1.1].

The proof for (b) is similar so we just give a sketch. Let G be the given p-group, and $\circ = \circ_G\colon G/Z(G) \times G/Z(G) \rightarrowtail G'$. Recall that we can compute generators for $A(\circ)$ as the solution of a system of equations. Hence, by Lemma 3.5, it suffices to find an idempotent $e \in A(\circ)$ such that $e^* = 1 - e$.

Using [W1, BW4, BW3] we begin by constructing the Jacobson radical, J, of A, and then decomposing A/J as a sum $S_1 \oplus \ldots \oplus S_m$ of $*$-simple (both simple and $*$-invariant) ideals. Each S_i is isomorphic to the adjoint ring of a nondegenerate alternating, symmetric, or Hermitian form (where in the Hermitian case we permit a degenerate field extension $K \oplus K$ – also called exchange); see [W5, Section 5].

In a $*$-simple ring, an idempotent e with $e^* = 1-e$ coincides with a decomposition of the associated form into a pair of totally isotropic subspaces, which are readily computed using a Gram-Schmidt type algorithm [W3]. Thus, within each S_i find an idempotent $\hat{e}_i$ with $\hat{e}_i^* = 1-\hat{e}_i$. Let $\hat{e} = \sum_i \hat{e}_i$ and use the idempotent lifting formula in [W1, Section 5.4] to lift $\hat{e} \in A/J$ to an idempotent $e \in A$ with $e^* = 1 - e$. □

We remark that one can lift idempotents more efficiently when p is odd by computing a $*$-invariant semisimple complement to the radical, thereby reducing the problem to the semisimple $*$-rings [BW4].

3.4. **The centrally indecomposable groups of genus 2.** We focus now on the centrally indecomposable groups of genus 2. Our immediate goal is to classify the adjoint rings of the commutation bimaps of such groups. The ultimate goal is to prove Theorem 1.2, but this must wait until Section 3.6.

We begin with a classification by Kronecker and Dieudonné [D] of pairs of matrices, which later led to classifications of pairs of forms by Scharlau [S1]. Independently – and prior to Scharlau – Bond [B3, p. 608] applied the same treatment to attempt to classify nilpotent Lie algebras of genus 2.

The following fundamental result is folklore (see, for example, [GG, Section 1]).

Lemma 3.7. *If $\{\Phi_1, \Phi_2\}$ is a pair of alternating forms on a finite-dimensional vector space V, then there is a decomposition $V = E \oplus F$ such that E and F are totally isotropic with respect to both forms.*

Let $\{\Psi_1, \Psi_2\}$ be a pair of $c \times d$ matrices with entries in a field k. As transformations from k^c to k^d we say that $\{\Psi_1, \Psi_2\}$ is *decomposable* if we can find bases for k^c and k^d with respect to which $\Psi_1 = \left[\begin{smallmatrix} \Psi_{11} & 0 \\ 0 & \Psi_{12} \end{smallmatrix}\right]$ and $\Psi_2 = \left[\begin{smallmatrix} \Psi_{21} & 0 \\ 0 & \Psi_{22} \end{smallmatrix}\right]$, and Ψ_{1j} has the same size as Ψ_{2j} for $j = 1, 2$. If no such bases exist then the pair is *indecomposable*. Indecomposable pairs are classified in the following classical result.

Theorem 3.8 (Kronecker-Dieudonné [D])**.** *If $\{\Psi_1, \Psi_2\}$ is an indecomposable pair of matrices with entries in a field k, then one of the following holds:*

(i) $\Psi_1, \Psi_2 \in \mathbb{M}_d(k)$ and there are bases such that $\Psi_1 = I_d$ and $\Psi_2 = C(a(x))$, where $a(x)$ is a power of an irreducible polynomial and $C(a(x))$ its companion matrix; or

(ii) $\Psi_1, \Psi_2 \in \mathbb{M}_{d,d+1}(k)$ and there are bases such that $\Psi_1 = [I_d|0]$ and $\Psi_2 = [0|I_d]$.

An algorithm for Theorem 3.8 is given in Section 5.1. The result asserts a canonical description of indecomposable pairs up to the action $\{\Psi_1, \Psi_2\} \mapsto \{X\Psi_1 Y, X\Psi_2 Y\}$. Note, constraining the problem so that $X = Y^{-1}$ (so the matrices are square) makes the classification problem wild. For, $x_i \mapsto \Psi_i$ defines a $k\langle x_1, x_2\rangle$-module on k^d, and conjugation of the pair $\{\Psi_1, \Psi_2\}$ by X is a module isomorphism; this is the definition of wild representations. Similarly, increasing from pairs of matrices to triples gives rise to another wild problem [BLS, BDL$^+$].

We use Theorem 3.8 now to classify pairs of forms associated to centrally indecomposable p-groups of genus 2.

Proposition 3.9. *If G is a centrally indecomposable p-group of genus 2 over a field k, then $\circ_G \colon k^d \times k^d \rightarrowtail k^2$ is isometric to a bimap represented by a pair*

$$\Phi_1 = \begin{bmatrix} 0 & \Psi_1 \\ -\Psi_1^{\mathrm{tr}} & 0 \end{bmatrix} \text{ and } \Phi_2 = \begin{bmatrix} 0 & \Psi_2 \\ -\Psi_2^{\mathrm{tr}} & 0 \end{bmatrix}, \tag{3.10}$$

of alternating forms, where the pair $\{\Psi_1, \Psi_2\}$ is given by Theorem 3.8 part (i) or (ii) according to whether d is even or odd, respectively.

Proof. Regard $V = G/Z(G)$ and $W = G'$ as k-spaces, so that $\dim_k W = 2$, and consider the k-bimap $\circ = \circ_G \colon V \times V \rightarrowtail W$. As in Section 2.1, $\circ$ is represented by a pair $\{\Phi_1, \Phi_2\}$ of alternating forms over k. As G is centrally indecomposable, by Lemma 3.5(i) $\{\Phi_1, \Phi_2\}$ is orthogonally indecomposable. Further, by Lemma 3.7, there is a decomposition $V = E \oplus F$ with $E \circ E = 0 = F \circ F$. Thus, the restriction of $\circ$ to $E \times F$ yields an indecomposable pair of matrices (the "corner blocks"). Using appropriate basis changes in E and F,

$$\Phi_1 = \begin{bmatrix} 0 & \Psi_1 \\ -\Psi_1^{\mathrm{tr}} & 0 \end{bmatrix} \text{ and } \Phi_2 = \begin{bmatrix} 0 & \Psi_2 \\ -\Psi_2^{\mathrm{tr}} & 0 \end{bmatrix},$$

where the pair $\{\Psi_1, \Psi_2\}$ is given by Theorem 3.8 part (i) or (ii) depending on whether $\dim_k V$ is even or odd, respectively. □

The dichotomy in Proposition 3.9 – and its eventual incarnation in Theorem 1.2 – is fundamental to our algorithm, and we introduce some helpful terminology from [BW3] for easy reference.

Definition 3.11. A centrally indecomposable group, G, of genus 2 is said to be *sloped* if it is type (i), and *flat* if it is type (ii). We extend the appropriate notion of sloped and flat to the associated k-bimap, $\circ_G$. (Note, if the k-dimension of $G/Z(G)$ is even then G is sloped, and otherwise it is flat.)

We stress that Proposition 3.9 is not a classification of centrally indecomposable groups of genus 2, even if their exponent is p. That would first require a classification of irreducible polynomials. Secondly – and much more troubling for our algorithms – *pairs of forms are not unique to a group of genus* 2.

Example 3.12. Let $k = \mathbb{Z}_3$, and put $a(x) = x^2 + 1$ and $b(x) = x^2 + x + 2$. The Heisenberg groups $H(k[x]/(a(x)^2))$ and $H(k[x]/(b(x)^2))$ are isomorphic (they are both over $\mathbb{F}_9$) and centrally indecomposable of genus 2, but there are sets of generators for these groups where the associated pairs of forms are non-isometric. E.g.:

$$\left\{ \begin{bmatrix} 0 & I_4 \\ -I_4 & 0 \end{bmatrix}, \begin{bmatrix} 0 & C(a(x)^2) \\ -C(a(x)^2)^{\mathrm{tr}} & 0 \end{bmatrix} \right\},$$
$$\left\{ \begin{bmatrix} 0 & I_4 \\ -I_4 & 0 \end{bmatrix}, \begin{bmatrix} 0 & C(b(x)^2) \\ -C(b(x)^2)^{\mathrm{tr}} & 0 \end{bmatrix} \right\},$$

and $C(a(x)^2)$ and $C(b(x)^2)$ are not conjugate.

Another way in which choice of generators removes a canonical relationship to Kronecker type arguments is seen by constructing groups of genus 2 as *quotients* of Heisenberg groups.

Example 3.13. Let $k = \mathbb{Z}_3$, and put $a_1(x) = x^4 + x^3 + x^2 + 1$, $a_2(x) = x^4 + 2x^2 + 2$, and $a_3(x) = x^4 + x^3 + 2x + 1$. Set $H_i = H(k[x]/(a_i(x)))$. As each $a_i(x)$ is irreducible $H_1 \cong H_2 \cong H_3$. Now, with respect to the natural basis $\{1, x, x^2, x^3\}$, define

$$M_i = \left\{ \begin{bmatrix} 1 & 0 & a + bx \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} : a, b \in k \right\} \leqslant H_i.$$

Evidently, $G_i = H_i/M_i$ has genus 2 over k and is centrally indecomposable. The assignment $x^2 \mapsto 2x^2 + x^3$ and $x^3 \mapsto x^2$ induces an isomorphism $G_2 \to G_3$. However, it can be shown (say, by applying the algorithm of Theorem 1.1) that $G_1 \ncong G_2$.

We will soon need the following consequence of Proposition 3.9.

Corollary 3.14. *Let G be a centrally indecomposable p-group of genus 2 over a field k, $A = A(\circ_G)$ its ring of adjoints, and $J = J(A)$ the Jacobson radical of A.*

(i) G is sloped if, and only if, $A/J \cong \mathbb{M}_2(K)$, where K/k a field extension and the induced involution on A/J is $\left[\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right] \mapsto \left[\begin{smallmatrix} d & -b \\ -c & a \end{smallmatrix}\right]$.

(ii) G is flat if, and only if, $A/J \cong K \oplus K$ with involution $(a, b) \mapsto (b, a)$.

Proof. By Lemma 3.5(i), the only idempotents of A with $e = e^*$ are 0 and 1. Furthermore, by Proposition 3.9, G is hyperbolic and so A has a hyperbolic idempotent $e^* = 1 - e$. As A is Artinian (in fact finite), idempotents lift over J, and by [W5, Section 5.4] they lift retaining the self-adjoint and hyperbolic relationships respectively. Therefore, A has such idempotents if, and only if, A/J has them. Now we apply a classification due to Osborn (see [W1, Theorem 4.26]) to assert the only choices for A/J are $M_2(K)$ and $K \oplus K$ together with the given involutions. In the first case the dimension of $G/Z(G)$ is even and hence corresponds to the sloped case. In the second case the group is flat. □

3.5. **Uniqueness of orthogonal and hyperbolic decompositions.** As we mentioned earlier, our algorithms for bimaps of genus 2 will utilize both types of decomposition described in the previous section. When we do so, we shall need to know that our particular choices are in fact generic. More precisely, we shall require the following transitivity facts. (Recall, from Proposition 2.2, that isoclinisms coincide with isomorphisms for groups of exponent p.)

Theorem 3.15. *If G is a finite p-group of genus 2 then the group of autoclinisms of G acts transitively on*

(a) the set of fully-refined central decompositions of G, and
(b) the set of hyperbolic pairs of G.

In fact the subgroup of autoclinisms that centralize $Z(G)$ is transitive on both sets.

Proof. For (a), refer to [W1, Theorem 6.6]. Corollary 3.14 tells us that $\circ_G$ has no indecomposable summands of orthogonal type, and so $\mathrm{Isom}(\circ_G)$ is transitive on its fully-refined orthogonal decompositions.

The proof of (b) is similar. By Witt's lemma, the isometries of a nondegenerate form act transitively on the set of hyperbolic bases. Then, using involutions, one lifts this action over the radical. □

We stress that central product decompositions *do not*, in general, possess such transitivity [W1, Theorem 1.1(ii)], so groups of genus 2 are somewhat special in this regard. Even so, Theorem 3.15 does not give rise to a theorem of Krull-Schmidt type [W1, Definition 2.6]. Indeed, as illustrated by the example below, identical sets of centrally indecomposable groups may occur as fully-refined central decompositions of non-isoclinic groups of genus 2. This hints at the difficulties in using central products within isomorphism tests.

Example 3.16. Let k be any field, and $K = k(\omega)$ a quadratic extension of k. Put $H = H_1(k) \times H_1(k) \times H_1(K)$, a direct product of Heisenberg groups, and let

$$N_1 = \left\langle \left(\begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, I_3, \begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \right), \left(I_3, \begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 0 & -\omega \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \right) \right\rangle$$

$$N_2 = \left\langle \left(\begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, I_3 \right), \left(I_3, \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \right) \right\rangle.$$

Then each N_i is normal in H, and the groups $G_i = H/N_i$ have genus 2 over k. Moreover, each group has a full-refined central decomposition consisting of two copies of a group X, and one copy of a group Y, yet G_1 and G_2 are non-isomorphic (in fact non-isoclinic). For, if ω has minimum polynomial $x^2 - ax - b$, then for $i = 1, 2$, the bimap $\circ_{G_i}$ is represented by $\left\{ \left[\begin{smallmatrix} 0 & I_4 \\ -I_4 & 0 \end{smallmatrix} \right], \left[\begin{smallmatrix} 0_4 & L_i \\ -L_i^t & 0_4 \end{smallmatrix} \right] \right\}$, where

$$L_1 = \begin{bmatrix} 0 & & & \\ & 1 & & \\ & & 0 & 1 \\ & & b & a \end{bmatrix} \qquad L_2 = \begin{bmatrix} 0 & & & \\ & 0 & & \\ & & 0 & 1 \\ & & b & a \end{bmatrix}.$$

Since L_1 has three distinct eigenvalues in K, and L_2 only two, the centralizers, $C(L_1)$ and $C(L_2)$, are non-isomorphic algebras. For $i = 1, 2$, by [BW3, Lemma

3.2], $A(\circ_{G_i})$ is isomorphic to $\mathbb{M}_2(C(L_i))$, so $A(\circ_{G_1})$ and $A(\circ_{G_2})$ are non-isomorphic algebras. It follows that G_1 and G_2 are non-isoclinic.

3.6. **A characterization of the indecomposable groups of genus 2.** We are almost ready to prove Theorem 1.2. Our approach requires that we examine the adjoint ring of the commutation bimap of these groups in greater depth. In particular, we provide a rather complete description of the bimap obtained by forming a tensor product over such rings (Theorem 3.18). As well as helping us prove Theorem 1.2, this will provide the foundation for the adjoint-tensor isomorphism test for the centrally indecomposable groups of genus 2 that we present in Section 5.

We will need the following convenient characterization of the adjoint ring of an indecomposable bimap of the sloped type.

Lemma 3.17 ([BW3, Lemma 3.2(i)])**.** *Let $\circ\colon k^d \times k^d \rightarrowtail k^2$ be an alternating bimap represented by a pair $\{\Phi_1, \Phi_2\}$ of forms with Φ_1 invertible. Then*

$$A(\circ) = C_{\mathbb{M}_d(k)}(\sigma), \qquad \text{where } \sigma = \Phi_2\Phi_1^{-1}.$$

In particular, $Z(A(\circ)) = k[\sigma] \cong k[x]/(m(x))$, where $m(x)$ is the minimum polynomial of σ. If $\circ$ is indecomposable, then $m(x) = a(x)^e$ with $a(x)$ irreducible.

Note, Lemma 3.17 requires only that Φ_1 is invertible and makes no assumption about the indecomposability of the bimap – we shall have more to say on this point in Remark 3.20 after we prove the following crucial result.

Theorem 3.18. *Let k be a field and $\circ\colon k^{2n} \times k^{2n} \rightarrowtail k^2$ an indecomposable, alternating bimap represented by a pair $\{\Phi_1, \Phi_2\}$ with Φ_1 invertible. Let $\sigma = \Phi_2\Phi_1^{-1}$, and let $m(x) \in k[x]$ be the minimal polynomial of σ. Then*

$$k^{2n} \otimes_{A(\circ)} k^{2n} = k^{2n} \wedge_{A(\circ)} k^{2n} \cong k[x]/(m(x)), \tag{3.19}$$

and $\otimes\colon k^{2n} \times k^{2n} \rightarrowtail k^{2n} \otimes_{A(\circ)} k^{2n}$ is a fully-nondegenerate alternating $k[x]/(m(x))$-form. Furthermore, the isomorphism in (3.19) can be computed in polynomial time.

Proof. By Lemma 3.7, there is a decomposition $k^{2n} = E \oplus F$ with $E \circ E = 0 = F \circ F$, and by Theorem 3.15 this decomposition is unique up to a choice of basis. Thus, we may assume

$$\Phi_1 = \begin{bmatrix} 0 & I_n \\ -I_n & 0 \end{bmatrix} \text{ and } \Phi_2 = \begin{bmatrix} 0 & \Psi \\ -\Psi^{\mathrm{tr}} & 0 \end{bmatrix},$$

where Ψ is in Rational Canonical Form, so that

$$\sigma = \Phi_2\Phi_1^{-1} = \begin{bmatrix} \Psi & 0 \\ 0 & \Psi^{\mathrm{tr}} \end{bmatrix}.$$

By Lemma 3.17, $A(\circ) = C_{\mathbb{M}_{2n}(k)}(\sigma) = \mathbb{M}_2(C_{\mathbb{M}_n(k)}(\Psi))$. The structure of centralizer matrices is well-studied and is determined by the representation of Ψ. In particular, there is a divisor chain $a_s(x)|a_{s-1}(x)|\cdots|a_1(x)$ of $m(x) = a_1(x)$ and, using companion matrices $C(a_i)$,

$$\Psi = \operatorname{diag}(C(a_1), \ldots, C(a_s)).$$

Correspondingly $E = E_1 \oplus \cdots \oplus E_s$ as a $k[x]$-module, with x acting as Ψ. The centralizer of Ψ is a chequered matrix, cf. [P, p. 42],

$$C_{\mathbb{M}_n(k)}(\Psi) = \left\{ [[M_{ij}]] : 1 \leqslant i, j \leqslant s,\ M_{ij} \in \operatorname{Hom}_{k[x]}(E_i, E_j) \right\}.$$

As E_1 is a faithful representation of $k[\Psi]$, so $\mathrm{Hom}_{k[x]}(E_1, E_i) = \mathrm{Hom}_{k[x]}(k[x], E_i) \cong E_i$ as $k[x]$-modules. Hence, there exists a matrix

$$M_\alpha = \begin{bmatrix} \tilde{\alpha}_1 & \tilde{\alpha}_2 & \cdots & \tilde{\alpha}_s \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{bmatrix} \in C_{\mathbb{M}_n(k)}(\Psi), \qquad e_1 M_\alpha = (\alpha_1, \ldots, \alpha_s) = \alpha \in E.$$

In particular, E and F are cyclic $k[x]$-modules, say

$$E = e_1 C_{\mathbb{M}_{d/2}(k)}(\Psi) \text{ and } F = f_1 C_{\mathbb{M}_{d/2}(k)}(\Psi).$$

We can now determine the structure of the bimap $\otimes_A$. Let $\alpha = (\alpha_1, \ldots, \alpha_s) \in E$ and $\beta = (\beta_1, \ldots, \beta_s) \in F$. If $\alpha \cdot \beta = \alpha_1\beta_1 + \cdots + \alpha_s\beta_s$ denotes the usual dot-product, then

$$\begin{aligned} \alpha \otimes_A \beta &= e_1 \begin{bmatrix} M_\alpha & \\ & 0_n \end{bmatrix} \otimes_A f_1 \begin{bmatrix} 0_n & \\ & M_\beta \end{bmatrix} \\ &= e_1 \begin{bmatrix} M_\alpha & \\ & 0_0 \end{bmatrix} \begin{bmatrix} M_\beta^{\mathrm{tr}} & \\ & 0_n \end{bmatrix} \otimes_A f_1 \\ &= (\alpha \cdot \beta)(e_1 \otimes_A f_1). \end{aligned}$$

In particular, $k^{2n} \otimes_A k^{2n} = k[x](e_1 \otimes f_1)$ is a cyclic $k[x]$-module, so the tensor product is a form. All of the necessary constructions are carried out in polynomial time so the result follows. □

Remark 3.20. Although our application to Theorem 1.2 concerns *indecomposable* bimaps of genus 2, Theorem 3.18 again requires only that Φ_1 is invertible (just like Lemma 3.17). This is explained in Section 6. We extend the notion of "sloped" to any alternating bimap of genus 2 represented by $\{\Phi_1, \Phi_2\}$ with Φ_1 invertible, and refer to $\sigma = \Phi_2\Phi_1^{-1}$ as a *slope* of $\circ$. The slope is crucial to the work in [BW3] but also features in earlier works such as [GG, BF].

We now can prove Theorem 1.2, which in our more general setting now says that if G is a centrally indecomposable p-group of genus 2 over a field k, then it is isoclinic to one of the following two types of groups:

(i) (sloped case) a quotient by a central subgroup N of a Heisenberg group,

$$H = \left\{ \begin{bmatrix} 1 & e & w \\ 0 & 1 & f \\ 0 & 0 & 1 \end{bmatrix} : \begin{array}{c} e, f, w \in k[x]/(a(x)^c), \\ a(x) \text{ irreducible} \end{array} \right\},$$

where as matrices $1 - N$ is a subspace of $1 - [H, H]$ of codimension 2; or

(ii) (flat case) the matrix group

$$H^\flat = \left\{ \left[\begin{array}{c|cccc|c} I_2 & e_1 & \cdots & e_m & 0 & w_1 \\ & 0 & e_1 & \cdots & e_m & w_2 \\ \hline & & & & & f_0 \\ & & I_{m+1} & & & \vdots \\ & & & & & f_m \\ \hline & & & & & 1 \end{array}\right] : e_i, f_j, w_\ell \in k \right\}.$$

Proof. Following Proposition 3.9, every centrally indecomposable group G of genus 2 determines a pair $\{\Phi_1, \Phi_2\}$ of alternating forms as in (3.10). It remains to connect the two possible matrix pairs to the corresponding matrix groups described in the theorem.

Suppose G is sloped and let $m(x) = a(x)^c$ be the minimum polynomial of $\circ_G : k^{2n} \times k^{2n} \rightarrowtail k^2$. Set $H = H(R)$, the Heisenberg group over $R = k[x]/(a(x)^c)$. Then the commutation bimap $\circ_H \colon R^2 \times R^2 \rightarrowtail R$ is an alternating R-*form*. Choose an isomorphism $\varphi : H/Z(H) \to G/Z(G)$ (both are isomorphic to k^{2n}).

By Theorem 3.18, $\circ_G$ factors through $\circ_H$ yielding a projection $\pi \colon R \to G'$ with $(u\varphi \circ_H v\varphi)\pi = u \circ_G v$. This gives rise to an isomorphism $\hat{\varphi} \colon H'/\ker\pi \to G'$, and $(\varphi, \hat{\varphi})$ is a pseudo-isometry from $\circ_G$ to $\circ_H$. Furthermore, if

$$N = \left\{ \begin{bmatrix} 1 & 0 & w \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} : w\pi = 0 \right\},$$

then G is isoclinic to H/N.

Next, we consider the flat case. Let E_{ij} indicate the matrix with 1 in position ij and 0 elsewhere. For $a \in \{1, \dots, m\}$, $b \in \{1, \dots, m+1\}$, $c \in \{1, 2\}$, and $t \in k$, set

$$\begin{aligned} M_a(t) &= I + tE_{1(a+2)} + tE_{2(a+3)}, \\ N_b(t) &= I + tE_{(2+b)(m+4)}, \text{ and} \\ Z_c(t) &= I + tE_{c(m+4)}. \end{aligned}$$

Then $H^\flat = \langle M_1(t), \dots, M_m(t), N_0(t), \dots, N_m(t), Z_1(t), Z_2(t) : t \in k \rangle$. From matrix multiplication we see $[M_a(t), M_{a'}(t')]$, $[N_b(t), N_{b'}(t')]$, $[M_a(t), Z_c(t')]$, $[N_b(t), Z_c(t')]$, and $[Z_c(t), Z_c'(t)]$ are all trivial. Furthermore,

$$[M_a(t), N_b(t')] = \begin{cases} Z_1(tt') & a = b, \\ Z_2(tt') & a + 1 = b, \\ 1 & \text{else.} \end{cases} \tag{3.21}$$

In particular, $[H^\flat, H^\flat] = \langle Z_1(t), Z_2(t) : t \in k \rangle$, and k is the centroid of the bimap of commutation in $H^\flat$. Also, $\{M_1(1), \dots, M_m(1), N_1(1), \dots, N_{m+1}(1)\}$ maps to a k-basis for $V = H^\flat/[H^\flat, H^\flat]$ and $\{Z_1(1), Z_2(1)\}$ is a k-basis for $W = [H^\flat, H^\flat]$. From (3.21), the structure constants coincide with a flat indecomposable pair of alternating forms. Thus, if G is centrally indecomposable and flat, then G and $H^\flat$ are isoclinic. □

3.7. **Generalized discriminants and Pfaffians.** We have previously stated that Theorem 1.2 is not enough to decide isomorphism among groups of genus 2 over a field k. So, what else is needed? The main result of this section provides a necessary and sufficient condition for isomorphism between groups of genus 2 whose indecomposable central factors are all sloped. This in turn gives rise to an isomorphism test that is effective when $|k|$ is small.

In the foregoing discussion of the bimap $\circ = \circ_G \colon V \times V \rightarrowtail W$ associated to such a group, G, we have worked exclusively with the k-space $V = G/Z(G)$. Now we turn our attention to $W = G' = k^2$. If $G' < Z(G)$ then $Z(G) = G' \times \mathbb{Z}_p^s$ and for some $G_0 < G$, $G = G_0 \times \mathbb{Z}_p^s$ and $G_0' = Z(G_0)$; thus, we may assume $G' = Z(G)$. If $\circ : G/G' \times G/G' \rightarrowtail G'$, then

$$1 \to \mathrm{Isom}(\circ) \to \Psi\mathrm{Isom}(\circ) \to \mathrm{Aut}(W)$$

is an exact sequence, and by Proposition 2.4,

$$1 \to C_{\mathrm{Aut}(G)}(W) \to \mathrm{Aut}(G) \to \mathrm{Aut}(W)$$

is an exact sequence. By assumption $\circ$ has centroid k a field. As $\mathrm{Aut}(G)$ acts on the centroid, $C(\circ) = k$, its action on W is k-semilinear so we may replace $\mathrm{Aut}(W)$ with $\Gamma\mathrm{L}(2,k)$. If $\{\Phi_1, \Phi_2\}$ is a pair of alternating forms representing $\circ$, we wish to study the action of the pseudo-isometry group $\Psi\mathrm{Isom}(\circ)$ on the 2-dimensional k-space spanned by this pair. In particular, we are interested in deciding when (and how) an element of $\Gamma\mathrm{L}(2,k)$ lifts to $\Psi\mathrm{Isom}(\circ)$. We begin by generalizing the notions of "discriminant" to arbitrary lists of square matrices, and of "Pfaffian" to pairs of alternating forms.

The discriminant of a bilinear form Ψ is the square class of the determinant. In this way it is invariant up to isometry, since $\det(X\Psi X^{\mathrm{tr}}) = \det(X)^2 \det(\Psi)$. For systems $\{\Psi_1, \ldots, \Psi_m\}$ of forms we define the *generalized discriminant* as follows:

$$\mathrm{disc}(\Psi_1, \ldots, \Psi_m) = \det(x_1\Psi_1 + \cdots + x_m\Psi_m) \in k[x_1, \ldots, x_m].$$

We shall work with such systems up to isotopism, which means we can modify by independent matrices X and Y to arrive at

$$\mathrm{disc}(X\Psi_1 Y, \ldots, X\Psi_m Y) = \det(X)\det(Y)\,\mathrm{disc}(\Psi_1, \ldots, \Psi_m).$$

Thus, $\mathrm{disc}(\Psi_1, \ldots, \Psi_m)$ is a homogenous polynomial defined only up to a non-zero scalar multiple – this is an interesting isotopism invariant so long as $m > 1$.

Next, consider a pair $\{\Phi_1, \Phi_2\}$ of alternating forms representing a sloped bimap of genus 2. There exist subspaces E and F of equal dimension relative to which,

$$(i = 1, 2) \qquad\qquad \Phi_i = \begin{bmatrix} 0 & \Psi_i \\ -\Psi_i^t & 0 \end{bmatrix},$$

so $\mathrm{disc}(\Phi_1, \Phi_2) = \mathrm{disc}(\Psi_1, \Psi_2)^2$. We therefore define the *generalized Pfaffian*,

$$\mathrm{Pf}(\Phi_1, \Phi_2) = \mathrm{disc}(\Psi_1, \Psi_2).$$

We will use a natural action of $\Gamma\mathrm{L}(2,k)$ on the homogeneous polynomials in $k[x,y]$. For $\hat{\alpha} = \left(\left[\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right], \tau\right) \in \mathrm{GL}(2,k) \rtimes \mathrm{Gal}(k)$, define

$$f^{\hat{\alpha}}(x,y) = f^{\tau}(ax + by, cx + dy).$$

We now integrate the Pfaffian of a sloped pair with our understanding of centrally indecomposable groups of genus 2 to interpret an isomorphism invariant introduced by Vishnevetskiĭ [V1, V2] in the case when $k = \mathbb{Z}_p$. A version of the following theorem was announced in [V1] but only the forward direction was proved. We need (a constructive version of) the converse for our isomorphism test so we provide a complete proof.

Theorem 3.22. *Let $\{\Phi_1, \Phi_2\}$ and $\{\Lambda_1, \Lambda_2\}$ be alternating k-forms, each written relative to a fully-refined orthogonal decomposition; so, for $i = 1, 2$,*

$$\Phi_i = \mathrm{diag}\left(\Phi_i^{(1)}, \ldots, \Phi_i^{(s)}\right) \qquad\qquad \Lambda_i = \mathrm{diag}\left(\Lambda_i^{(1)}, \ldots, \Lambda_i^{(t)}\right).$$

For $\hat{\alpha} \in \Gamma\mathrm{L}(2,k)$, there is a pseudo-isometry $(\alpha, \hat{\alpha})$ from $\{\Phi_1, \Phi_2\}$ to $\{\Lambda_1, \Lambda_2\}$ if, and only if, $s = t$ and there is a permutation σ of $\{1, \ldots, s\}$ such that for all i,

$$\mathrm{Pf}\left(\Phi_1^{(i)}, \Phi_2^{(i)}\right)^{\hat{\alpha}} \equiv \mathrm{Pf}\left(\Lambda_1^{(i\sigma)}, \Lambda_2^{(i\sigma)}\right) \pmod{k^\times}. \tag{3.23}$$

Proof. We begin with the forward direction. Assume $(\alpha, \hat{\alpha}) \in \Gamma\mathrm{L}(2n, k) \times \Gamma\mathrm{L}(2, k)$ is a pseudo-isometry from $\{\Phi_1, \Phi_2\}$ to $\{\Lambda_1, \Lambda_2\}$. The transitivity result in Theorem 3.15 may be recast in the language of orthogonal decompositions for the associated bimaps. In particular, there is an isometry carrying the basis of the fully-refined orthogonal decomposition of $\{\alpha\Phi_1\alpha^{\mathrm{tr}}, \alpha\Phi_2\alpha^{\mathrm{tr}}\}$ to that of $\{\Lambda_1, \Lambda_2\}$. As the former has s terms, and the latter t terms, it follows that $s = t$.

Let $\overline{\Phi}_i = \alpha\Phi_i\alpha^{\mathrm{tr}}$ (the semilinear action is coordinatewise) and so

$$\overline{\Phi}_i = \mathrm{diag}\left(\overline{\Phi}_i^{(1)}, \ldots, \overline{\Phi}_i^{(s)}\right)$$

is a fully-refined orthogonal decomposition of $\left\{\overline{\Phi}_1, \overline{\Phi}_2\right\}$. Since $(\alpha, \hat{\alpha})$ is a pseudo-isometry, it follows that there exists a permutation σ of $\{1, \ldots, s\}$ such that

$$\left(\overline{\Phi}_1^{(i)}, \overline{\Phi}_2^{(i)}\right)^{\hat{\alpha}} = \left(\Lambda_1^{(i\sigma)}, \Lambda_2^{(i\sigma)}\right).$$

Hence, observing that

$$\mathrm{Pf}\left(\Phi_1^{(i)}, \Phi_2^{(i)}\right)^{\hat{\alpha}} = \mathrm{Pf}\left(\left(\Phi_1^{(i)}, \Phi_2^{(i)}\right)^{\hat{\alpha}}\right),$$

and that $\mathrm{Pf}\left(\Phi_1, \Phi_2\right) \equiv \mathrm{Pf}\left(\overline{\Phi}_1, \overline{\Phi}_2\right) \pmod{k^\times}$, we see that (3.23) holds for this σ.

Conversely, suppose $\hat{\alpha} = (\hat{\mu}, \tau) \in \mathrm{GL}(2, k) \rtimes \mathrm{Gal}(k)$ satisfies (3.23). It suffices to find, for each i, an independent α_i such that $(\alpha_i, \hat{\alpha})$ is a pseudo-isometry from $\left\{\Phi_1^{(i)}, \Phi_2^{(i)}\right\}$ to $\left\{\Lambda_1^{(i\sigma)}, \Lambda_2^{(i\sigma)}\right\}$. Indeed, if this is the case, then $\big((\alpha_1 \oplus \cdots \oplus \alpha_s)\Sigma(\sigma^{-1}), \hat{\alpha}\big)$ is a pseudo-isometry from $\{\Phi_1, \Phi_2\}$ to $\{\Lambda_1, \Lambda_2\}$, where $\Sigma(\sigma^{-1})$ is the permutation matrix associated to σ^{-1}. Therefore, we assume $\circ$ and $\bullet$ (represented by $\{\Phi_1, \Phi_2\}$ and $\{\Lambda_1, \Lambda_2\}$, respectively) are indecomposable. By Lemma 2.3, a semilinear pseudo isometry must use the same Galois automorphism τ in the domain and codomain. Hence, we are only concerned with finding μ such that $\big((\mu, \tau), (\hat{\mu}, \tau)\big)$ is a pseudo-isometry from $\{\Phi_1, \Phi_2\}$ to $\{\Lambda_1, \Lambda_2\}$. We consider two special cases for $\hat{\mu}$ before we treat the general case.

First, suppose that $\hat{\mu}$ fixes an indeterminant of $k[x, y]$, modulo $k^\times$. By Proposition 3.9, we may assume that there are bases relative to which

$$\mathrm{Pf}(\Phi_1, \Phi_2) = \det(xI + yC) \qquad \mathrm{Pf}(\Lambda_1, \Lambda_2) = \det(xI + yD).$$

Then, either $\mathrm{Pf}(\Phi_1, \Phi_2)^{\hat{\alpha}} \equiv \det(xI + yM)$ for some M, or $\mathrm{Pf}(\Phi_1, \Phi_2)^{\hat{\alpha}} \equiv \det(xN + yC)$ for some N. Note that either option is equivalent to $\mathrm{Pf}(\Lambda_1, \Lambda_2) = \det(Ix + Dy)$ since $\mathrm{Pf}(\Phi_1, \Phi_2)^{\hat{\alpha}} \equiv \mathrm{Pf}(\Lambda_1, \Lambda_2)$. As $\{\Phi_1, \Phi_2\}$ and $\{\Lambda_1, \Lambda_2\}$ represent indecomposable bimaps, $\{I, M\}$ (or $\{N, C\}$ as the case may be) and $\{I, D\}$ are indecomposable pairs of matrices. Hence, by the Kronecker-Dieudonné theorem, there are matrices X and Y such that $X\{I, D\}Y = \{I, C\}^\tau = \{I, C^\tau\}$. If $\mu = \mathrm{diag}(X, Y^{\mathrm{tr}})$ and $\alpha = (\mu, \tau)$, then $(\alpha, \hat{\alpha})$ is a pseudo-isometry from $\{\Phi_1, \Phi_2\}$ to $\{\Lambda_1, \Lambda_2\}$, so any lower or upper triangular $\hat{\mu}$ can be lifted.

Next, suppose $\hat{\mu}$ interchanges x and y. Thus, we may assume there are bases such that

$$\mathrm{Pf}(\Phi_1, \Phi_2) = (xI + yC) \qquad \mathrm{Pf}(\Lambda_1, \Lambda_2) = (xD + yI).$$

Arguing as before, there are matrices X and Y such that $X\{D, I\}Y = \{I, C^\tau\}$ and if $\mu = \mathrm{diag}(X, Y^{\mathrm{tr}})$ and $\alpha = (\mu, \tau)$, then $(\alpha, \hat{\alpha})$ is a pseudo-isometry from $\{\Phi_1, \Phi_2\}$ to $\{\Lambda_1, \Lambda_2\}$.

In the general case, $\hat{\mu}$ is the product of $\hat{\beta}\hat{\gamma}\hat{\delta}$ where $\hat{\beta}$ and $\hat{\gamma}$ fix x or y (modulo $k^\times$), and $\hat{\delta}$ transposes or fixes them (an LUP-decomposition). Here, we lift $\hat{\mu}$ with three iterations using the two special cases already treated. Therefore, with $\alpha = (\mu, \tau)$, the pair $(\alpha, \hat{\alpha})$ is a pseudo-isometry from $\{\Phi_1, \Phi_2\}$ to $\{\Lambda_1, \Lambda_2\}$, and so the theorem follows. □

From Theorem 3.22, since the equivalence is modulo $k^\times$, we need only consider $\hat{\alpha} \in \mathrm{P\Gamma L}(2, k)$. Hence, we obtain the following corollary.

Corollary 3.24. *Let $\{\Phi_1, \Phi_2\}$ and $\{\Lambda_1, \Lambda_2\}$ be pairs of sloped $d \times d$ forms over $\mathbb{F}_q$, for $q = p^e$. There exists an algorithm that determines if the two are pseudo-isometric using $O(q^3 e + d^3)$ field operations.*

Proof. The rational canonical form of a $d \times d$ matrix is computed using $O(d^3)$ field operations [S4]. □

4. The Adjoint-Tensor Method

Having developed the necessary foundation, we turn now to our isomorphism tests. To emphasize that our algorithms apply only to finite groups and fields we shall henceforth write $\mathbb{F}_q$ in place of k, where $\mathbb{F}_q$ is an extension of $\mathbb{Z}_p$. Via Proposition 2.4 questions of isomorphism between finite p-groups of class 2 are reduced to ones of pseudo-isometry between $\mathbb{F}_q$-bimaps. Details of this reduction – and the isomorphism tests it leads to – are given in Section 7. Our current focus is the following problem.

PseudoIsometry ($\circ$, $\bullet$)
Given: alternating $\mathbb{F}_q$-bimaps $\circ, \bullet\colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^e$
Return: a pseudo-isometry from $\circ$ to $\bullet$, if such exists.

We can, in principle, return *all* pseudo-isometries from $\circ$ to $\bullet$ as a coset of the group $\Psi\mathrm{Isom}(\circ)$. That group is often the focus of attention because it relates directly to the automorphism group of a p-group. We shall concentrate here on testing for pseudo-isometry and explain how to adapt our methods to finding generators for $\Psi\mathrm{Isom}(\circ)$ in Section 6.4.

The first simplification is to reduce from general ($\mathbb{F}_q$-semilinear) pseudo-isometries to $\mathbb{F}_q$-linear pseudo-isometries. The following result states that testing for $\mathbb{F}_q$-linear pseudo-isometry is the heart of matter.

Theorem 4.1. *Given an algorithm to solve $\mathbb{F}_q$-linear pseudo-isometry that runs in time $t(n)$, $n = d \log q$, then there is an algorithm that solves $\mathbb{F}_q$-semilinear pseudo-isometry in time $O(t(n) \log q)$.*

Proof. As the actions on radicals have no restrictions, we may assume that the given bimaps are fully-nondegenerate. Let $\circ, \bullet\colon V \times V \rightarrowtail W$ be an instance of $\mathbb{F}_q$-semilinear pseudo-isometry. For each $\sigma \in \mathrm{Gal}(\mathbb{F}_q)$, proceed as follows. Define

$$u * v = u^\sigma \circ v^\sigma, \qquad\qquad u \# v = (u \bullet v)^\sigma,$$

and test if there is a $\mathbb{F}_q$-linear pseudo-isometry $(\varphi, \hat{\varphi})$ from $*$ to $\#$. If so, it follows that $((\varphi, \sigma), (\sigma, \hat{\varphi})) \in (\mathrm{GL}(V) \ltimes \mathrm{Gal}(\mathbb{F}_q)) \times (\mathrm{Gal}(\mathbb{F}_q) \rtimes \mathrm{GL}(W))$ is a pseudo-isometry from $\circ$ to $\bullet$, namely $u(\varphi, \sigma) \circ v(\varphi, \sigma) = u\varphi * v\varphi = (u \circ v)(\sigma, \hat{\varphi})$. If this fails for every $\sigma \in \mathrm{Gal}(\mathbb{F}_q)$ then $\circ$ and $\bullet$ are not pseudo-isometric. □

In view of Theorem 4.1, we henceforth assume that all pseudo-isometries are linear over the field of definition. In particular, $\Psi\mathrm{Isom}(\circ\colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^e)$ *will mean the group of* $\mathbb{F}_q$*-linear pseudo-isometries of* $\circ$*.*

The question of testing alternating bimaps for pseudo-isometry is one that arises also in the generic method for group isomorphism – the *p-group generation algorithm* – though framed in cosmetically different terms; see [O]. The basic approach is as follows. As both bimaps are alternating, they factor through the alternating tensor bimap $\wedge\colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^d \wedge \mathbb{F}_q^d$ with induced maps $\hat{\circ}, \hat{\bullet}\colon \mathbb{F}_q^d \wedge \mathbb{F}_q^d \to \mathbb{F}_q^e$. The group $\mathrm{GL}(d, \mathbb{F}_q)$ acts naturally on $\mathbb{F}_q^d \wedge \mathbb{F}_q^d$, and $\circ$ and $\bullet$ are pseudo-isometric if, and only if, an element of $\mathrm{GL}(d, \mathbb{F}_q)$ maps $\ker \hat{\circ}$ to $\ker \hat{\bullet}$. Thus, to determine pseudo-isometry, we must solve a *subspace transporter problem*, which is notoriously difficult even for "well understood" actions like the exterior square representation. In practice, it is possible to proceed by a direct orbit calculation only for quite modest values of d and q.

If $\circ$ and $\bullet$ have a constrained structure – such as the bimaps arising in [LW] – specialized techniques may be developed to compute orbits efficiently. Inspired by the need to bridge the gap between slow, generic methods, and very fast, highly specialized ones, in [BW2] the first and third authors proposed a new general technique called the *adjoint-tensor method*. The method, which we outline in general below, is particularly well-suited to the alternating bimaps of genus 2; most of remaining content of the paper is concerned with the application of adjoint-tensor to this case.

PseudoIsometry ($\circ, \bullet\colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^e$)

1. Compute $A = A(\circ)$ and $A(\bullet)$.
2. Test if there exists $\rho \in \mathrm{GL}(d, \mathbb{F}_q)$ with $A(\bullet)^\rho = A$; if not, return `false`.
   Replace $\bullet$ with an isometric bimap, $\star$, so that $A = A(\circ) = A(\star)$.
3. Compute the kernels of the induced maps $\hat{\circ}, \hat{\star}\colon \mathbb{F}_q^d \wedge \mathbb{F}_q^d \to \mathbb{F}_q^2$.
4. Construct generators for $\Psi\mathrm{Isom}(\wedge_A)$.
5. Find $(\varphi, \hat{\varphi}) \in \Psi\mathrm{Isom}(\wedge_A)$ with $(\ker \hat{\circ})\hat{\varphi} = \ker \hat{\star}$;
   return the pseudo-isometry $(\rho\varphi, \hat{\varphi})$ from $\circ \to \bullet$.

FIGURE 4.1. The adjoint-tensor approach to solving PseudoIsometry.

Step 1 computes the two adjoint rings, which is no worse than solving a system of ed^2 equations in $2d^2$ variables. In certain situations – notably sloped genus 2 bimaps – one can extend the practical range by avoiding these large linear systems [BW3].

No polynomial-time solution is known for the general problem in step 2. It asks whether subalgebras of $\mathbb{M}_d(\mathbb{F}_q)$ are conjugate which, beyond just being isomorphic, requires that they are identically represented on $\mathbb{F}_q^d$. This leads to the notion of *module similarity*, a problem which was shown in [BW1] to be as hard as graph isomorphism.

To understand step 3, recall from Section 3.3 that the adjoint ring $A = A(\circ)$ is the largest ring, B, such that $\circ$ factors through the tensor product $\mathbb{F}_q^d \otimes_B \mathbb{F}_q^d$. Since, for us, the bimap $\circ$ is alternating, it additionally factors through the exterior product $\mathbb{F}_q^d \wedge_A \mathbb{F}_q^d$ (cf. Theorem 3.18), so there is an induced map $\hat{\circ}\colon \mathbb{F}_q^d \wedge_A \mathbb{F}_q^d \to \mathbb{F}_q^e$

such that

$$(\forall u, v \in V) \qquad\qquad u \circ v = (u \wedge v)^{\hat{\circ}}.$$

Computing $\ker \hat{\circ}$ and $\ker \hat{\star}$ amounts to solving a system of $O(d^2)$ linear equations.

Step 4 builds the group that acts on $V \wedge_A V$. As we noted earlier, $\mathrm{GL}(V)$ acts naturally on the components of the traditional exterior square $V \wedge V$. To respect the tensor over A we must instead use the group $\Psi\mathrm{Isom}(\wedge_A)$, the structure of which is described in [BW2, Theorem 4.5]. The description requires one to compute the normalizer of A, and the complexity of this problem depends critically on structural properties of A and on its representation.

The final component (step 5) is the same as the conclusion of the p-group generation algorithm described earlier. Once again $\circ$ and $\star$ are pseudo-isometric if, and only if, $\ker \hat{\circ}$ and $\ker \hat{\star}$ are in the same orbit, this time under the action of $\Psi\mathrm{Isom}(\wedge_A)$. Hence, we must solve the subspace transporter problem for the representation of $\Psi\mathrm{Isom}(\wedge_A)$ on $\mathbb{F}_q^d \wedge_A \mathbb{F}_q^d$.

In summary, steps 2 (module similarity), 4 (normalizers of matrix rings), and 5 (subspace transporters) are each known to be at least as hard as graph isomorphism [BW1, Theorem 1.2; BW2, Theorem 4.5(iii); LM]. It seems, then, that adjoint-tensor merely turns one difficult problem into three! The idea, though, is that each new "hard problem" is either smaller in size, or has a controlled structure that admits a more efficient solution. *This is exactly the case for groups of genus* 2.

5. Indecomposable Bimaps of Genus 2

We now restrict to bimaps of genus 2 and develop an effective algorithm for PseudoIsometry in this case. We start in this section by further restricting to (orthogonally) indecomposable bimaps. Our goal is the following result.

Theorem 5.1. *There is a polynomial-time algorithm that, given indecomposable, alternating bimaps $\circ, \bullet \colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$ of genus* 2*, decides if the bimaps are pseudo-isometric and, if so, constructs a k-linear pseudo-isometry, namely $(\varphi, \hat{\varphi}) \in \mathrm{GL}(d, \mathbb{F}_q) \times \mathrm{GL}(2, \mathbb{F}_q)$ such that $u\varphi \bullet v\varphi = (u \circ v)\hat{\varphi}$ for all $u, v \in \mathbb{F}_q^d$. The algorithm is deterministic if p is bounded and Las Vegas otherwise.*

Recall, the qualification of Las Vegas versus deterministic in Theorem 5.1 (and in later theorems) arises only from the need to factor polynomials.

5.1. **Standard indecomposable pairs of matrices.** To prove Theorem 5.1 we apply the "flat-sloped" dichotomy of Theorem 1.2 to the associated pairs $\{\Phi_1, \Phi_2\}$ of alternating forms. It will be helpful to select a basis relative to which

$$\Phi_1 = \begin{bmatrix} 0 & \Psi_1 \\ -\Psi_1^{\mathrm{tr}} & 0 \end{bmatrix} \text{ and } \Phi_2 = \begin{bmatrix} 0 & \Psi_2 \\ -\Psi_2^{\mathrm{tr}} & 0 \end{bmatrix}, \tag{5.2}$$

and $\{\Psi_1, \Psi_2\}$ is given by the appropriate part of Theorem 3.8. We begin by finding a totally isotropic decomposition for the pair (see Lemma 3.7). This is done in polynomial time by finding a hyperbolic pair of idempotents in the adjoint ring of the pair, as we did in the proof of Theorem 3.6(ii). By changing to a basis that respects this totally isotropic decomposition, we obtain a pair of forms as in (5.2) with $\{\Psi_1, \Psi_2\}$ an arbitrary indecomposable pair of matrices. It remains to find matrices X, Y such that $\{X\Psi_1 Y, X\Psi_2 Y\}$ has the desired form.

The sloped case, namely Theorem 3.8(i), has been discussed from the point of view of algorithms in several recent papers; see [GG, BW3] for example. The conversion depends only on the sloped aspect, and so works for decomposable pairs.

Suppose $\Psi_1, \Psi_2 \in \mathbb{M}_{n,n+1}(\mathbb{F}_q)$ is an indecomposable pair, where $d = 2n+1$. Compute $X \in \mathbb{M}_n(\mathbb{F}_q), Y \in \mathbb{M}_{n+1}(\mathbb{F}_q)$ such that $X\Psi_1 Y = [I_n|0]$, the standard matrix for Ψ_1, using Gaussian elimination. We now modify X and Y so that $X\Psi_1 Y = [I_n|0]$ and $X\Psi_2 Y = [0|I_n]$. We do this by successive approximations.

First, write $X\Psi_2 Y = [U|u^{\mathrm{tr}}]$, and find $B \in \mathrm{GL}_n(\mathbb{F}_q)$ such that $BUB^{-1} = R$ is in generalized Jordan normal form. Reassign $X := BX$ and $Y := Y \left[\begin{smallmatrix} B^{-1} & 0 \\ 0 & 1 \end{smallmatrix}\right]$. As the pair is indecomposable, R is a single companion matrix, say $R = \left[\begin{smallmatrix} 0 & I_{n-1} \\ \alpha & v' \end{smallmatrix}\right]$. Secondly, write $X\Psi_2 Y = [R|v^{\mathrm{tr}}]$ and find T in the cyclic algebra generated by R sending v^{tr} to $(0 \ldots 01)^{\mathrm{tr}}$. Reassign $X := TX$. Finally, write $X\Psi_2 Y = \left[\begin{smallmatrix} 0 & I_{n-1} & 0 \\ \beta & b' & 1 \end{smallmatrix}\right]$, put $b := (\beta\ b') \in \mathbb{F}_q^n$, and reassign $Y := Y \left[\begin{smallmatrix} I_n & 0 \\ -b & 1 \end{smallmatrix}\right]$.

5.2. **The flat case.** It is immediate from our discussion of this case in the preceding section that two flat, indecomposable bimaps of genus 2 are isometric, and Theorem 5.1 holds in this case. Recall, however, we shall eventually require generators for $\Psi\mathrm{Isom}(\circ)$. We address this problem for general bimaps later in Section 6.4, but we can resolve the matter now for flat, indecomposable bimaps of genus 2.

Proposition 5.3. *If $\circ\colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$ is a flat, indecomposable bimap of genus 2, then there is an epimorphism $\Psi\mathrm{Isom}(\circ) \to \Gamma\mathrm{L}(2, \mathbb{F}_q)$ with kernel $\mathrm{Isom}(\circ)$.*

Proof. For $e = (e_1, \ldots, e_n) \in \mathbb{F}_q^n$, define $M = M(e) = \left[\begin{smallmatrix} e_1 & \cdots & e_n & 0 \\ 0 & e_1 & \cdots & e_n \end{smallmatrix}\right]$ and set

$$E = \{M(e)\colon e \in \mathbb{F}_q^n\} \leqslant \mathbb{M}_{2\times(n+1)}(\mathbb{F}_q).$$

Then, the usual matrix multiplication

$$\times\colon E \times \mathbb{M}_{(n+1)\times 1}(\mathbb{F}_q) \rightarrowtail \mathbb{M}_{2\times 1}(\mathbb{F}_q)$$

is described by the system of forms $\Psi = \{[I_n|0], [0|I_n]\}$. As $\circ$ is given by a pair of forms $\Phi_i = \left[\begin{smallmatrix} 0 & \Psi_i \\ -\Psi_i^{tr} & 0 \end{smallmatrix}\right]$, it follows that every isotopism $(\alpha, \beta; \gamma)$ of $\times$ induces a pseudo-isometry $(\alpha \oplus \beta, \gamma)$ of $\circ$, so it suffices to show that $\Gamma\mathrm{L}(2, \mathbb{F}_q)$ lifts to autotopisms of $\times$ acting faithfully on $\mathbb{M}_{2\times(n+1)}(\mathbb{F}_q)$.

For $e = (e_1, \ldots, e_n) \in \mathbb{F}_q^n$, define $f_e(x, y) = e_1 x^n + \cdots + e_i x^{n-i} y^i + \cdots + e_n y^n \in \mathbb{F}_q[x, y]$. Then $M_e \mapsto f_e(x, y)$ is a linear bijection from $E = \{M(e)\colon e \in \mathbb{F}_q^n\}$ to the set of homogeneous polynomials in $\mathbb{F}_q[x, y]$ of degree n. Let $\rho\colon \Gamma\mathrm{L}(2, \mathbb{F}_q) \to \Gamma\mathrm{L}(n+1, \mathbb{F}_q)$ denote the faithful representation arising from the natural action of $\Gamma\mathrm{L}(2, \mathbb{F}_q)$ on the latter. For $g \in \Gamma\mathrm{L}(2, \mathbb{F}_q)$, define

$$M(e)\lambda_g := gM(e)(g^{-1}\rho).$$

Then $(\lambda_g, g\rho; g)$ is an isotopism of $\bullet$, and the result follows. □

5.3. **The sloped case.** Recall that an alternating bimap $\circ\colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$ of genus 2 is sloped if we can represent it by a pair $\{\Phi_1, \Phi_2\}$ with Φ_1 nondegenerate. Our goal is to complete the proof of Theorem 5.1 by presenting a test for k-linear pseudo-isometry between two sloped, indecomposable bimaps $\circ, \bullet\colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$.

We will use the adjoint-tensor method of Section 4, referring to the pseudo-code in Figure 4.1. Recall that we must resolve three problems:

Line 2 Given adjoint algebras $A(\circ)$ and $A(\bullet)$, find $\rho \in \mathrm{GL}(d, \mathbb{F}_q)$ with $A(\bullet)^\rho = A(\circ)$ (if such exists).
Line 4 Given $A = A(\circ)$, build generators for $\Psi\mathrm{Isom}(\wedge_A)$.
Line 5 Solve the "transporter problem": given subspaces U, V of $\mathbb{F}_q^d \otimes_A \mathbb{F}_q^d$ find $(\varphi, \hat{\varphi}) \in \Psi\mathrm{Isom}(\wedge_A)$ sending U to V, or prove that no such $(\varphi, \hat{\varphi})$ exists.

We consider each problem in turn.

5.3.1. *Conjugating the adjoint algebras.* As we noted in Section 4, conjugacy of algebras is very hard in general, but an efficient solution exists in our setting. This relies on the special nature of adjoint algebras for sloped bimaps of genus 2. Any such bimap $\circ$ is represented by a pair $\{\Phi_1, \Phi_2\}$ of alternating forms with Φ_1 invertible, and its slope $\sigma = \Phi_2\Phi_1^{-1}$ is invariant under basis change in $\mathbb{F}_q^d$ – that is, invariant in $\mathrm{Isom}(\circ)$. By Lemma 3.17, $A(\circ) = C_{\mathbb{M}_d(\mathbb{F}_q)}(\sigma)$, so

$$Z(A(\circ)) = \mathbb{F}_q[\sigma] \cong \mathbb{F}_q[x]/(m(x)),$$

where $m(x)$ is the minimum polynomial of σ. The conjugacy problem for cyclic algebras has an efficient solution.

Theorem 5.4 ([BW1, Theorem 1.3])**.** *There is a polynomial-time algorithm that, given cyclic algebras $A = \mathbb{F}_q[\alpha]$, $B = \mathbb{F}_q[\beta]$, for $\alpha, \beta \in \mathbb{M}_d(\mathbb{F}_q)$, finds $\rho \in \mathrm{GL}(d, \mathbb{F}_q)$ with $A^\rho = B$, or decides that no such ρ exists.*

This leads to a resolution of our first problem.

Corollary 5.5. *There is a polynomial-time algorithm that, given sloped bimaps $\circ, \bullet\colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$ of genus 2, finds $\rho \in \mathrm{GL}(d, \mathbb{F}_q)$ such that $A(\circ)^\rho = A(\bullet)$, or decides that no such ρ exists.*

Proof. By Lemma 3.17, $A(\circ)$ and $A(\bullet)$ are centralizers of slopes $\sigma_\circ$ and $\sigma_\bullet$, respectively. Furthermore, $A(\circ)$ and $A(\bullet)$ are conjugate if, and only if, their centers are conjugate. The result now follows from Theorem 5.4. □

5.3.2. *The properties of $\wedge_A$.* To describe $\Psi\mathrm{Isom}(\wedge_A)$ we need the following result.

Theorem 5.6 ([BW2, Theorem 1.5])**.** *If $\circ\colon V \times V \rightarrowtail W$ is an alternating k-bimap with adjoint ring $A = A(\circ)$, then $\Psi\mathrm{Isom}_k(\wedge_A)$ is faithfully represented on V as*

$$N^*(A) = \{g \in \mathrm{GL}(d, k)\colon A^g = A \text{ and } (x^g)^* = (x^*)^g \text{ for all } x \in A\}.$$

Using the general structure of $N^*(A)$ laid out in [BW2, Theorem 4.5] together with Theorem 3.18, we gain a very detailed understanding of $\Psi\mathrm{Isom}(\wedge_A)$ when $\circ\colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$ is a sloped, indecomposable bimap of genus 2. Put

$$K = \mathbb{F}_q[x]/(m(x)),$$

where $m(x)$, the minimal polynomial of the slope of $\circ$, is a power of an irreducible polynomial. Hence, $K \cong L[t]/(t^e)$ where $L/\mathbb{F}_q$ is an algebraic field extension, and

$$1 \longrightarrow \mathrm{Isom}(\wedge_A) \longrightarrow \Psi\mathrm{Isom}(\wedge_A) \longrightarrow \Gamma\mathrm{L}(1, K) \longrightarrow 1$$

is a short exact sequence, where $\mathrm{Isom}(\wedge_A)$ is the kernel of the action of $\Psi\mathrm{Isom}(\wedge_A)$ on $V \wedge_A V$. The algorithms to find generators and further structure of isometry groups were given in [BW4, BW3]. Hence, the group we must understand is

$$\Psi\mathrm{Isom}(\wedge_A)/\mathrm{Isom}(\wedge_A) \cong \Gamma\mathrm{L}_{\mathbb{F}_q}(1, K) = K^\times \rtimes \mathrm{Aut}_{\mathbb{F}_q}(K).$$

The group $\mathrm{Aut}_{\mathbb{F}_q}(K) = \mathrm{Aut}_{\mathbb{F}_q}(L[t]/(t^e))$ satisfies

$$1 \to \Sigma \to \mathrm{Aut}_{\mathbb{F}_q}(K) \to \Gamma\mathrm{L}_{\mathbb{F}_q}(1, L) \to 1,$$

where $\Sigma = C_{\mathrm{Aut}(K)}((t)/(t^2))$ is a quotient of the *Nottingham group*, a well-studied pro p-group [LGM, Section 12.4]. Generators for $\Gamma\mathrm{L}_{\mathbb{F}_q}(1, L) = L^\times \rtimes \mathrm{Gal}(L/\mathbb{F}_q)$ are known, and Σ consists of *substitution automorphisms*,

$$\Lambda_{a(t)} \colon p(t) \mapsto p(a(t)),$$

where $a(t) = t + a_2t^2 + \cdots$; it is generated by $\{\Lambda_{t+t^2}, \Lambda_{t+t^3}\}$.

5.3.3. *Solving the transporter problem.* Our final concern is to solve the transporter problem: *given subspaces U, V of $\mathbb{F}_q^d \wedge_A \mathbb{F}_q^d$, find $(\alpha, \hat{\alpha}) \in \Psi\mathrm{Isom}(\wedge_A)$ sending U to V, or prove that no such $(\alpha, \hat{\alpha})$ exists.*

Our algorithm handles the Galois group of L by exhaustion. We work with the remaining part, namely with $G := \Sigma K^\times$, in a more refined manner. To facilitate our computations, we choose generators for G that produce a convenient factorization. First, as a consequence of Wedderburn's Principal Theorem, $K^\times$ factorizes as $Q_1 \rtimes G_1$, with Q_1 unipotent, and G_1 isomorphic to the multiplicative group of a field. Secondly, as we saw above, there is an analogous factorization, $Q_2 \rtimes G_2$, of Σ. Put $Q := Q_1Q_2$, and $J := J(K)$, the Jacobson radical of K. The crucial properties of the factorization QG_1G_2 for our purpose are as follows:

(i) Q is a unipotent group;
(ii) there are fields L_1, L_2 such that $G_1 = L_1^\times$ and $G_2 = L_2^\times$; and
(iii) G_1 acts faithfully on the k-space K/J, and G_2 acts faithfully on J/J^2.

Before proceeding further, we require two different "transporter" algorithms that will solve our problem in special cases. The proof of the following result generalizes an earlier algorithm of L. Rónyai developed for the case of fields; see [LW, Lemma 4.8].

Lemma 5.7. *Let R be a subalgebra of $\mathbb{M}_d(\mathbb{F}_q)$. Given $\mathbb{F}_q$-subspaces X, Y of R with $\dim X = \dim Y > 0$, in polynomial time one can find $r \in R^\times$ with $Xr = Y$, or decide that no such r exists.*

Proof. Let $t = \dim X = \dim Y > 0$. Find a basis for the $\mathbb{F}_q$-space

$$\mathcal{S} = \{a \in R \colon Xa \subseteq Y\}$$

as follows. Let $b_1, \ldots, b_n$ be an $\mathbb{F}_q$-basis for R. Fix bases for X and Y. Let $y_1, \ldots, y_t$ be the basis for Y, and write $y_q = \sum_{p=1}^n \gamma_{pq} b_p$ for $1 \leqslant q \leqslant t$. Now, for each basis element $x = \sum_{i=1}^n \alpha_i b_i$ of X, we want all scalars $z_1, \ldots, z_n, w_1, \ldots, w_t$ such that

$$\left(\sum_{i=1}^n \alpha_i b_i\right)\left(\sum_{j=1}^n z_j b_j\right) = \sum_{i,j=1}^n \alpha_i z_j b_i b_j = \sum_{p=1}^t \sum_{q=1}^n \gamma_{pq} w_p b_q.$$

Writing each b_ib_j as a linear combination of $b_1, \ldots, b_n$ (these are the *structure constants* of R relative to our chosen basis), a basis for $\mathcal{S}$ is obtained as the solution of the resulting linear system in the unknowns $z_1, \ldots, z_n, w_1, \ldots, w_t$ by projecting onto the z_i coordinates.

Evidently, if $\mathcal{S} = 0$, no $r \in R$ exists with $Xr \subseteq Y$, so we may assume that $\mathcal{S} \neq 0$. As we require a unit of R transporting X to Y, we must locate an injective element

of $\mathcal{S}$ if such exists. We present a deterministic method, but in practice such an element is found more efficiently by random search.

Compute $\mathcal{T} = \{b \in R \colon Yb \subseteq X\}$ as above (interchanging the roles of X and Y). Form the set $\mathcal{ST} = \{st : s \in \mathcal{S}, t \in \mathcal{T}\} \subset \mathrm{End}(X)$. Then using the algorithm of [BL, Theorem 2.4] we prove that there are no invertible elements in $\mathcal{ST}$, or we construct an invertible element z of the subring generated by $\mathcal{ST}$ as a product $z = s_1t_1s_2t_2 \cdots s_nt_n$, $s_i \in \mathcal{S}$, $t_i \in \mathcal{T}$. In the latter case, s_1 is injective. □

Remark 5.8. We intend to apply Lemma 5.7 in the case when X and Y have codimension 2 in R. By translating the problem to the dual space of R, we can solve the transporter problem instead for spaces of dimension 2. This reduces the complexity of computing the $\mathbb{F}_q$-spaces $\mathcal{S}$ and $\mathcal{T}$ by a factor of $O(d)$.

The following is a special case of the deeper theorem of [L2, Theorem 3.2(7)]. It is also known by many as the "unipotent stabilizer algorithm" (see [S2], for example).

Lemma 5.9. *Let Q be a unipotent subgroup of $\mathrm{GL}(d, \mathbb{F}_q)$. Given subspaces X, Y of $\mathbb{F}_q^d$, in polynomial time one can find $u \in Q$ with $Xu = Y$ if such u exists.*

We can now complete the description of our algorithm. Recall that U and V are given $\mathbb{F}_q$-subspaces of K, J is the Jacobson radical of K, and $G = L_1^\times L_2^\times Q$. We wish to decide if there exists $g \in G$ such that $Ug = V$.

First, construct the representation of L_1 on K/J, and use Lemma 5.7 to find $g_1 \in L_1^\times$ such that $Ug_1 \equiv V \pmod{J}$, if such exists. Put $U_1 = Ug_1$. Next, construct the representation of L_2 on J/J^2, and use Lemma 5.7 again to find $g_2 \in L_2^\times$ such that $U_1Jg_2 \equiv VJ \pmod{J^2}$, if such exists. Put $U_2 = U_1g_2$. Finally, use Lemma 5.9 to find $w \in Q$ with $U_2w = V$, if such exists. Return $g := g_1g_2w$. Note, if we failed to construct any one of the elements g_1, g_2, w, then there is no $g \in G$ transporting U to V.

Proof of Theorem 5.1. The correctness of the algorithms presented in Sections 5.2 and 5.3 has already been established. It remains to analyze complexity.

The sloped case in Section 5.3 requires more analysis, and we proceed one subsection at a time. First, conjugating the algebra $A(\bullet)$ to $A(\circ)$ is done by Corollary 5.5. Secondly, building the tensor product $\mathbb{F}_q^d \wedge_A \mathbb{F}_q^d$ and generators of $\Psi\mathrm{Isom}(\wedge_A)$ is done in polynomial time in Section 5.3.2. That leaves Section 5.3.3, which requires more care. For each $\gamma \in \mathrm{Gal}(L/\mathbb{F}_q)$ we seek $g \in \mathrm{Aut}_{\mathbb{F}_q}(K)$ with $(V\gamma)g = U$. This uses two calls to Lemma 5.7, and one call to Lemma 5.9, which are both polynomial time. The overall complexity is therefore polynomial, since $|\mathrm{Gal}(L/\mathbb{F}_q)| \leqslant \frac{d}{2}$. □

6. General Bimaps of Genus 2

We now consider arbitrary alternating bimaps $\circ, \bullet \colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$ of genus 2. Much of the work has already been done in the indecomposable setting above, but we must now combine the results of various indecomposables. Here, the theory becomes difficult. As Example 3.16 shows, for instance, indecomposable factors may be glued together in different ways to produce bimaps that are not pseudo-isometric. In spite of these challenges we prove the following extension of Theorem 5.1.

Theorem 6.1. *There is an algorithm that, given alternating bimaps $\circ, \bullet \colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$ of genus 2, constructs $(\varphi, \hat{\varphi}) \in GL(d, \mathbb{F}_q) \times \mathrm{GL}(2, \mathbb{F}_q)$ such that $u\varphi \bullet v\varphi = (u \circ v)\hat{\varphi}$*

for all $u, v \in \mathbb{F}_q^d$, or proves no such pair exists. The algorithm is polynomial time if q is bounded, or if the number of pairwise pseudo-isometric indecomposable summands of the input bimaps is bounded. If p is bounded the algorithms are deterministic, otherwise they are Las Vegas.

Using Theorem 3.6(i), one first constructs a fully-refined orthogonal decomposition of the input bimaps. By Theorem 3.15(i), the multiset of terms in such a decomposition is unique up to pseudo-isometry. Hence, if the terms in the two decompositions cannot be paired up pseudo-isometrically, then the bimaps themselves are not pseudo-isometric. In particular, if the multisets of dimensions of indecomposables are different for the two bimaps, then they are not pseudo-isometric. Furthermore, assuming the dimensions of the flat indecomposables are compatible, $\circ$ and $\bullet$ are pseudo-isometric if, and only if, their restrictions to the sum of the sloped parts are pseudo-isometric. Hence, we may assume that the indecomposable factors of each bimap are sloped.

We reiterate that deciding pseudo-isometry of $\circ$ and $\bullet$ is not as straight-forward as matching up isomorphic sloped indecomposable factors – more subtlety is required. We present two rather different approaches. The first is very effective when $|\mathbb{F}_q|$ is small, and is based directly on the theory developed in Section 3.7. The second, which we use for larger fields, is the adjoint-tensor method. Before proceeding we must first address a curiosity that can arise in our new setting.

6.1. **A rare configuration.** Recall that $\circ, \bullet \colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$ are nondegenerate bimaps whose indecomposable summands are sloped. We would like $\circ$ and $\bullet$ to be sloped *globally*, meaning that each may be represented by a pair of forms $\{\Phi_1, \Phi_2\}$ with Φ_1 nondegenerate. Certainly, an initial choice of basis can produce a representative pair of degenerate forms, as the following example for $d = 4$ shows:

$$\Phi_1 = \left[\begin{array}{cc|cc} & & 1 & \\ & & & 0 \\ \hline -1 & & & \\ & 0 & & \end{array}\right] \qquad \Phi_2 = \left[\begin{array}{cc|cc} & & 0 & \\ & & & 1 \\ \hline 0 & & & \\ & -1 & & \end{array}\right],$$

Here, though, the associated bimap is pseudo-isometric to a bimap represented by $\{\Phi_1 + \Phi_2, \Phi_2\}$, and $\Phi_1 + \Phi_2$ is nondegenerate. When the field is sufficiently large, we can always make such adjustments. In particular, the following holds.

Lemma 6.2. *Let k be an infinite field, and $\circ \colon k^d \times k^d \rightarrowtail k^2$ a nondegenerate, alternating k-bimap of genus 2. Then $\circ$ is sloped if, and only if, all of its indecomposable factors are sloped.*

Proof. The forward direction is clear. For the converse, suppose

$$(i = 1, 2) \qquad \Phi_i = \operatorname{diag}(\Phi_i^{(1)}, \ldots, \Phi_i^{(t)}),$$

represents $\circ$ and respects a fully-refined orthogonal decomposition, where each $\{\Phi_1^{(j)}, \Phi_2^{(j)}\}$ is sloped. A linear combination of Φ_1, Φ_2 is nondegenerate if, and only if, some evaluation of $\operatorname{disc}(\Phi_1, \Phi_2) \in k[x, y]$ does not vanish. By assumption, each $\operatorname{disc}(\Phi_1^{(i)}, \Phi_2^{(i)}) \neq 0$ (as a polynomial), so $\operatorname{disc}(\Phi_1, \Phi_2) = \prod_i \operatorname{disc}(\Phi_1^{(i)}, \Phi_2^{(i)}) \neq 0$. As k is infinite, there is a point not on the variety of $\operatorname{disc}(\Phi_1, \Phi_2)$. □

For finite fields, the situation is more delicate.

Lemma 6.3. *For every finite field $\mathbb{F}_q$, there is an integer d and an alternating bimap $\circ\colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$, all of whose indecomposable summands are sloped, such that every pair $\{\Phi_1, \Phi_2\}$ of forms representing $\circ$ consists of degenerate matrices.*

Proof. Consider a pair $\{\Phi_1, \Phi_2\}$ representing an alternating bimap $\circ\colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$. If no nondegenerate linear combination of Φ_1, Φ_2 exists, then the Pfaffian $\operatorname{Pf}(\Phi_1, \Phi_2)$ vanishes on all of $PG(1, \mathbb{F}_q)$. This means that $\prod_{\omega \in \mathbb{F}_q}(x - \omega y) \in \mathbb{F}_q[x, y]$ divides $\operatorname{Pf}(\Phi_1, \Phi_2)$. For each $\omega \in \mathbb{F}_q$,

$$\operatorname{Pf}\left(\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}, \begin{bmatrix} 0 & \omega \\ -\omega & 0 \end{bmatrix}\right) = x - \omega y.$$

The orthogonal sum of all such pairs yields a pair of forms whose discriminant vanishes on $PG(1, \mathbb{F}_q)$, but whose indecomposable summands are all sloped. □

Fortunately, our analysis comes to the rescue. The following scholium allows us to treat $\mathbb{F}_q$ as a "small" field whenever such a configuration occurs.

Lemma 6.4. *If $\circ\colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$ is an alternating non-sloped bimap, all of whose indecomposable summands are sloped, then $q < d$.*

Proof. Let $\{\Phi_1, \Phi_2\}$ represent $\circ\colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$. Then $\operatorname{Pf}(\Phi_1, \Phi_2)$ is a homogeneous polynomial of degree $\frac{d}{2}$. As $\circ$ is non-sloped, however, the proof of Lemma 6.3 shows that $\operatorname{Pf}(\Phi_1, \Phi_2)$ is divisible by $\prod_{\omega \in \mathbb{F}_q}(x - \omega y)$ of degree q. Hence, $q \leqslant \frac{d}{2} < d$. □

6.2. **Pfaffian test for small fields.** Let $\circ, \bullet\colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$ be two given bimaps whose indecomposable summands are all sloped. Write each bimap relative to a fully-refined orthogonal decomposition, and represented, as in Theorem 3.22, by a pair $\{\Phi_1, \Phi_2\}$ with $\Phi_i = \operatorname{diag}\left(\Phi_i^{(1)}, \ldots, \Phi_i^{(s)}\right)$.

To each pair we associate a collection $\{\operatorname{Pf}(\Phi_1^{(i)}, \Phi_2^{(i)})\colon 1 \leqslant i \leqslant s\}$ of homogeneous polynomials. Then, using Theorem 3.22, we can test whether or not $\circ$ and $\bullet$ are pseudo-isometric by exhaustively checking every element $\hat{\alpha} \in \operatorname{P\Gamma L}(2, \mathbb{F}_q)$ to see if it yields an equivalence between the two collections. Moreover, Theorem 3.22 is constructive: for suitable $\hat{\alpha}$ we can compute $\alpha \in \operatorname{\Gamma L}(d, \mathbb{F}_q)$ such that $(\alpha, \hat{\alpha})$ is a pseudo-isometry from $\circ$ to $\bullet$.

The complexity of the algorithm outlined above contains an unavoidable factor of $|\operatorname{P\Gamma L}(2, \mathbb{F}_q)|$ for the exhaustive search, cf. Corollary 3.24. In practice it works well when q is small. Recall, if $\{\Phi_1, \Phi_2\}$ satisfies the hypotheses of Lemma 6.4, then $q < d$, and we regard d as small. In particular, for the remainder of this section we assume that alternating bimaps of genus 2 over large fields are sloped.

6.3. **Adjoint-tensor test for large fields.** The shortcut isomorphism test described in the preceding section, while very effective in practical settings, has an unavoidable factor of $O(q^3 e d^3)$ in its complexity. Hence, the performance of this technique deteriorates quickly as the size of q increases. We therefore adapted the adjoint-tensor method to this more general setting. To illustrate the improvement, for primes p increasing from 3 to 257 we generated five pairs of isomorphic groups of order p^{10} and compared the performance of the Pfaffian method to that of the adjoint-tensor method. The results are displayed in Figure 6.1, where the plots indicate runtimes for each individual isomorphism test. Unsurprisingly, the timing for the Pfaffian method is inconsistent as its search through $\operatorname{P\Gamma L}(2, p)$ varies from group to group.

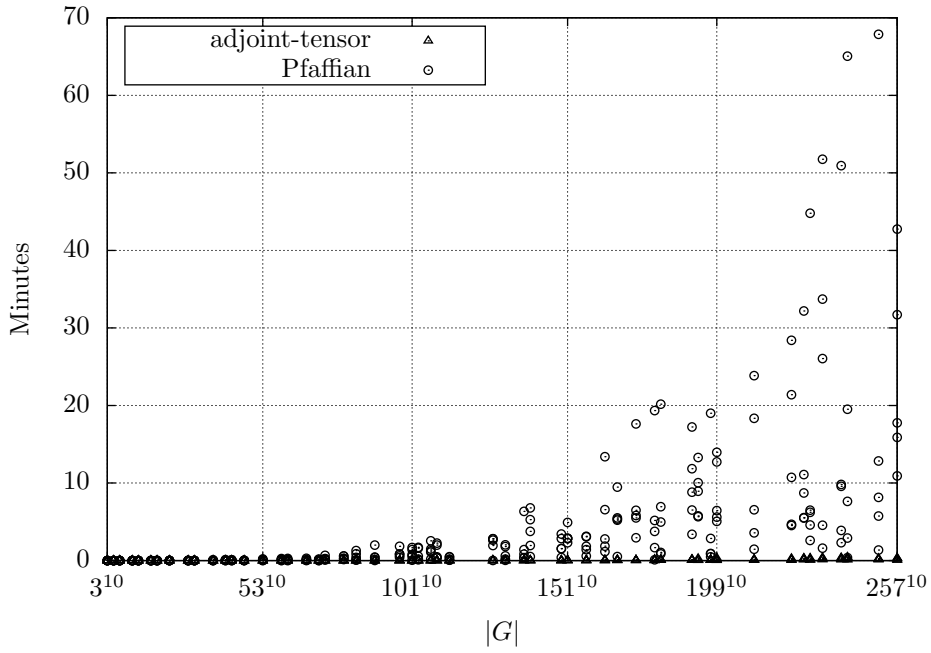

FIGURE 6.1. Comparison of adjoint-tensor method with Pfaffian method for random p-groups of genus 2 as we let p increase.

The algorithm proceeds exactly as described in Section 5.3. Difficulties arise because the structure of $N^*(A)$ is more complex than in the indecomposable case, and because we handle that additional structure by brute force. This gives rise to the rather less elegant complexity statement in Theorem 6.1. We now discuss all of the subtleties that arise in moving from indecomposable bimaps to general bimaps, and indicate how we handle them.

The first subtlety occurs, in fact, prior to the main algorithm. Recall, we have assumed that the indecomposable summands are all sloped. Our presentation of the adjoint-tensor method presumes that the given bimaps are sloped "globally". This is the importance of Lemma 6.4: if this happens not to be the case, then q is small relative to d and we use Section 6.2 instead.

The crucial point suggested above is that, unlike the approach in Section 6.2, we treat the input bimaps globally (rather than working with indecomposable summands). The adjoint algebras $A(\circ)$ and $A(\bullet)$ are still centralizers of single matrices, and hence the conjugacy problem in Section 5.3.1 goes through unchanged.

Moving on to the tensor product $T = \mathbb{F}_q^d \wedge_A \mathbb{F}_q^d$, once again there is little new; Theorem 3.18 (as noted in Remark 3.20) applies in this more general setting. That is, if $\sigma = \Phi_2\Phi_1^{-1}$ is a slope of $\circ$, and $m(x)$ its minimal polynomial, then $T \cong K = \mathbb{F}_q[x]/(m(x))$ as $\mathbb{F}_q$-modules. The only difference is that now K may have multiple primary components (it's not usually local) but we can compute independently within each primary component.

We turn now to the structure of $\Psi\mathrm{Isom}(\wedge_A) = N^*(A)$. We refer, once again, to the general structure theorem in [BW2, Theorem 4.5].

First, note that it's possible to have a subgroup of permutation matrices inside $N^*(A)$ arising from the representation of A. More precisely, $N^*(A)$ permutes the isoptypic components of the decomposition of K into primary components. Recall that primary components V_i $(i = 1, 2)$ are *isotypic* if they have minimal polynomial p_i^n, p_i irreducible, with $\deg p_1 = \deg p_2$, and where the V_i have identical Jordan block structures. We denote this permutation subgroup of $N^*(A)$ by Π. It is possible, provided $q \geqslant \frac{d}{2}$, for $|\Pi|$ to be as large as $\left(\frac{d}{2}\right)!$.

Secondly, $K/J(K)$ is a product of fields (as opposed to a single field). Therefore, the subgroup Γ of $\Gamma\mathrm{L}(1, K)$, which was previously a single Galois group, may contain a direct product of Galois groups. Hence, $|\Gamma|$ may be as large as $2^{\frac{d}{4}}$.

We now turn to the final step of the algorithm in Section 5.3.3. We proceed exactly as before, but instead of looping over Γ, we now loop over $\Gamma\Pi$. Observe that both $|\Gamma|$ and $|\Pi|$ are bounded under the additional hypotheses of the last assertion in Theorem 6.1, which is what yields polynomial time in that case.

6.4. **The group of pseudo-isometries of a bimap of genus 2.** Recall that our test for pseudo-isometry between given bimaps $\circ, \bullet \colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$ promises the set of *all* such pseudo-isometries (if such is needed). It does so by additionally returning generators for the group $\Psi\mathrm{Isom}(\circ)$.

Recall, in view of Theorem 4.1, we consider just $\mathbb{F}_q$-linear pseudo-isometries.

PseudoIsometryGroup ($\circ$)

Given: an alternating $\mathbb{F}_q$-bimap $\circ \colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^e$.
Return: (generators for) the group $\Psi\mathrm{Isom}(\circ)$.

Again, we focus on genus 2, and consider first the situation where q is considered small, as in Section 6.2. Here, there is very little to be said. We proceed – as though testing for pseudo-isometry between $\circ$ and itself – by listing all $\hat{\varphi} \in \mathrm{GL}(2, \mathbb{F}_q)$ and testing whether $\hat{\varphi}$ lifts to a pseudo-isometry $(\varphi, \hat{\varphi})$ of $\circ$. When we have exhausted the elements of $\mathrm{GL}(2, \mathbb{F}_q)$ we have the entire group $\Psi\mathrm{Isom}(\circ)$.

Next, suppose that q is large. Any $\mathbb{F}_q$-linear pseudo-isometry of $\circ$ preserves a basic decomposition of $\circ$ into its flat and sloped parts. We saw in Proposition 5.3 that the pseudo-isometry group of the flat part induces the full $\mathrm{GL}(2, \mathbb{F}_q)$ on $\mathbb{F}_q^2$, and this result is constructive in that it provides a lift of any given $\hat{\varphi} \in \mathrm{GL}(2, \mathbb{F}_q)$ to a pseudo-isometry of the flat part. Thus, in view of Section 6.1, it suffices to construct $\Psi\mathrm{Isom}(\circ)$ when $\circ \colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$ is sloped.

This, in fact, is somewhat easier than deciding pseudo-isometry because we need not concern ourselves with conjugating adjoint algebras. In fact, referring to the pseudo-code in Section 4, everything remains the same in an algorithm for PseudoIsometryGroup until Line 5. Here, instead of seeking a single element $(\varphi, \hat{\varphi}) \in \Psi\mathrm{Isom}(\wedge_A)$ mapping $\ker \hat{\circ}$ to $\ker \hat{\star}$, we require the full stabilizer in $\Psi\mathrm{Isom}(\wedge_A)$ of $\ker \hat{\circ}$, say. One solves such "stabilizer" problems using exactly the same machinery we used for the "transporter" problems at no additional cost.

In sum, we have proved the following.

Theorem 6.5. *There is a deterministic algorithm that, given an alternating $\mathbb{F}_q$-bimap $\circ \colon \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^2$ of genus 2, constructs generators for $\Psi\mathrm{Isom}(\circ)$. The algorithm is polynomial time if either q is bounded, or if the number of pairwise pseudo-isometric indecomposable summands of the input bimaps is bounded.*

7. Proof of Theorem 1.1

In fact we prove a stronger version of our main Theorem 1.1, which includes groups that are direct products of groups of genus at most 2.

Part (a). First we must recognize when a group G is a direct product of groups of genus 2 (possibly for multiple fields and characteristics). The algorithms of [W4, Section 5] may be used to write $G = G_1 \times \cdots \times G_s$ with each G_i directly indecomposable. Hence, by the Krull-Remak-Schmidt theorem, we may assume G_1 and G_2 are directly indecomposable groups.

Recall from Section 2.3 that we assume our computational model allows us to decide if a group G is a p-group of class 2 and, if so, to construct its associated bimap $\circ_G \colon V \times V \rightarrowtail W$. Recall, also, that the centroid $C(\circ)$ can be computed by solving a system of linear equations. Using [BO, Section 2.2], compute a Wedderburn decomposition $C(\circ) = K \oplus J(C(\circ))$. As the centroid of a directly indecomposable p-group of class 2 is local, it follows that K is a field. Since $C(\circ)$ is commutative, K is unique, and the K-dimension of W is well-defined. Deciding whether G has genus at most 2 is now a simple check whether $\dim_K W \leqslant 2$.

Part (b.1). Suppose we are given directly indecomposable groups G_1 and G_2 of genus at most 2 over fields K_i; we must determine the set of isoclinisms $G_1 \to G_2$. Note, the set is empty if the G_i have different genera, or if $K_1 \ncong K_2$ (which can be decided using [BW1, Lemma 3.5]). By fixing coordinates we may assume the bimaps of commutation are represented on common vectors spaces: we have alternating, fully-nondegenerate bimaps $\circ_1, \circ_2 : \mathbb{F}_q^d \times \mathbb{F}_q^d \rightarrowtail \mathbb{F}_q^e$ for some $1 \leqslant e \leqslant 2$.

Now, if $e = 1$, then the $\circ_i$ are alternating nondegenerate forms over a local ring, and have a symplectic basis that is unique up to pseudo-isometry (see, for example, [MH, Chapter I, Corollary 3.5]). Using a Gram-Schmidt process (see [W3], for example) construct a pseudo-isometry $(\varphi, \hat{\varphi}) \colon \circ_1 \to \circ_2$. The group of pseudo-isometries of $\circ_i$ is $\Gamma\mathrm{Sp}(d, \mathbb{F}_q)$, and standard generators for these groups are well known. Return the set of isoclinisms $\circ_1 \to \circ_2$ as the coset $\Gamma\mathrm{Sp}(d, \mathbb{F}_q)(\varphi, \hat{\varphi})$.

Otherwise, $e = 2$. Here, we use our algorithms for Theorems 4.1 and 6.1 to find a pseudo-isometry $(\varphi, \hat{\varphi}) : \circ_1 \to \circ_2$ if one exists. If none exist, return the empty set. Else, construct generators for $\Psi\mathrm{Isom}(\circ_1)$ using Theorem 6.5 and extend to semilinear pseudo-isometries using Theorem 4.1. Return $\Psi\mathrm{Isom}(\circ_1)(\varphi, \hat{\varphi})$.

That the coset of pseudo-isometries $\circ_1 \to \circ_2$ corresponds to the coset of isoclinisms $G_1 \to G_2$ follows from Theorem 2.1 (Baer Correspondence).

Part (b.2). To complete the proof of Theorem 1.1, we must upgrade isoclinism to isomorphism in the cases when the input groups have exponent $p > 2$. Note, first, that we can recognize when a given p-group, G, has exponent p by computing a matrix, P, representing the linear map $xZ(G) \mapsto x^p$: then G has exponent p if, and only if, $P = 0$.

Given an isoclinism $(\varphi, \hat{\varphi})$, there is an induced isomorphism $\Phi : G_1 \to G_2$; for explicit construction see [W1, Proposition 3.10]. It remains to compute generators for the full automorphism group of G_1 and return $\mathrm{Aut}(G_1)\Phi$. Note, $\mathrm{Aut}(G_1) \cong \mathrm{Hom}_{\mathbb{F}_q}(\mathbb{F}_q^d, \mathbb{F}_q^2) \rtimes \Psi\mathrm{Isom}(\circ_1)$ where, for $(\varphi, \hat{\varphi}) \in \Psi\mathrm{Isom}(\circ_1)$ and $\mu \in \mathrm{Hom}_{\mathbb{F}_q}(\mathbb{F}_q^d, \mathbb{F}_q^2)$,

$$\mu^{(\varphi, \hat{\varphi})} = \varphi^{-1} \mu \hat{\varphi}.$$

The representation of $(\mu, (\varphi, \hat{\varphi}))$ on G is given in the same manner as [W1, Proposition 3.10(ii)].

Complexity. The k-linear pseudo-isometry problem dominates the complexity of the procedure, and this is estimated in Theorem 6.1. It remains to cast this estimate in terms of the input groups, where $|G_i'| = \mathbb{F}_q^2$. First, the Pfaffian test, which runs deterministically in time $O(q^3 \log q + d^{2\omega} \log q)$, translates as $O(|G_i'|^{3/2} \log |G_i'| + (\log |G_i|)^{2\omega}))$. For the adjoint-tensor estimate in Theorem 6.1, observe that t – an upper bound on the number of pairwise non-pseudo-isometric orthogonal factors of the given bimaps – translates to a bound on the number of pairwise nonisomorphic central factors of the given groups. The corresponding complexity in that case is $O(t! + (\log |G_i|)^{2\omega})$, and the proof of Theorem 1.1 is now complete. □

8. Implementation and Performance

As mentioned in Section 1.1, we have implemented the algorithms presented in Sections 4–7 in the computer algebra system MAGMA. Our implementation, which is available upon request, makes essential use of the StarAlgebra package implemented by the first and third authors [BW4]. Although there are areas where performance can be improved, the plots in Figures 1.1 and 6.1 illustrate the efficacy of our implementation. All tests were carried out on an Intel® Xeon® E5-1620, 3.60 GHz microprocessor, running MAGMA V2.21-11.

We now comment on the results depicted in Figure 1.1 (henceforth referred to as Experiment A), and on further experiments designed to probe the behavior of the implementation in different circumstances.

Experiment A. We constructed random pairs of groups of genus 2 as follows. For fixed d we generated a pair $\{\Phi_1, \Phi_2\}$ of skew-symmetric $d \times d$ matrices with entries in $\mathbb{F}_5$. Next, we built a 5-group, G, of genus 2 as a PC-group with commutator relations determined by the entries of $\{\Phi_1, \Phi_2\}$. We then chose random $g \in \mathrm{GL}(d, \mathbb{F}_5)$ and $h \in \mathrm{GL}(2, \mathbb{F}_5)$, computed $\{\Psi_1, \Psi_2\} = \{g\Phi_1 g^{\mathrm{tr}}, g\Phi_2 g^{\mathrm{tr}}\}^h$, and used these matrices to define another PC-group, H which, by construction, is isomorphic to G. Finally, we used our implementation to test for isomorphism between G and H. The test was repeated 10 times for each even d between 4 and 254, and 3 times for each odd d between 3 and 199, for a total 1557 tests. For 16 groups it took longer than 70 minutes to construct an isomorphism; the most extreme example was a centrally indecomposable group of order 5^{256} which took a little over two hours. By construction, all groups constructed for odd d were flat and for even d were sloped.

We have shown that, for bounded primes, the asymptotic complexity of our algorithms is $O(d^{2\omega})$, which matches the complexity of solving systems of linear equations in $(1 + o(1))d^2$ equations and variables – we refer to this as "d^2 linear algebra". To see this behavior in our implementation we looked at the ratio $\log t_{\cong}(d, p)/\log t_{LA}(d, p)$, where $t_{\cong}(d, p)$ is the time to test isomorphism of p-groups of genus 2 and order p^{d+2}, and $t_{LA}(d, p)$ is the corresponding time to solve a random linear system of $(1 + o(1))d^2$ equations in $(1 + o(1))d^2$ variables over a field of size p (in particular we varied the number of variables compared to equations to sample between situations that are both under- and over-determined). The results are reported in Figure 8.1 and they demonstrate that the performance of our implementation is aligned with theory. Evidently, when applied to sloped instances, the complexity converges more rapidly to that of d^2 linear algebra than it does when

applied to flat instances, but Figure 8.1 shows that both cases are trending in the right direction.

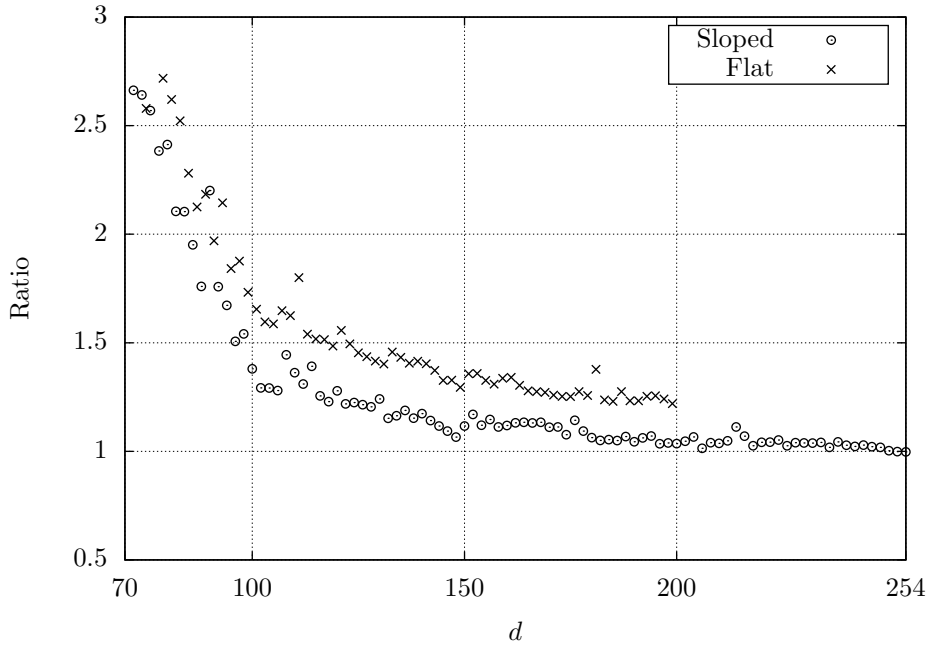

FIGURE 8.1. Graph showing the ratio of the logarithm of our runtimes against the logarithm of the time needed to solve $(1+o(1))d^2$ linear systems as $d = \log_p |G| - 2$ increases.

The variance in runtimes in Figure 1.1 for the sloped case (in contrast to the smooth graph for the flat groups) is due to the construction and manipulation of the adjoint algebra. In the case of a sloped group, G, one constructs $A(\circ_G)$ very quickly using the methods of [BW3]. For these groups, the completion time is affected if the Jacobson radical of $A(\circ_G)$ is nontrivial, or if the natural module for $A(\circ_G)$ decomposes into many blocks. Information about the number of blocks and sizes of the largest blocks for the groups in our experiment are given in Figures 8.2 and 8.3, respectively. As one can see, for sloped groups there is considerable variability in the block structure. With flat groups, on the other hand, there is usually just a single indecomposable block, and the runtime is always dominated by the construction of the adjoint algebra. This accounts for the relatively smooth graph for flat groups and the spiky graph for sloped groups.

Experiment B. We conducted an experiment to examine the behavior of our implementation when given pairs of input groups that are unlikely to be isomorphic. We fixed $p = 1021$, selected random even integers, d, between 20 and 40 and, as in Experiment A, built p-groups, G, of genus 2 with $d = \log_p |G| - 2$ from a random pair of alternating forms. For each G we then constructed a group H from an independent pair of alternating forms of the same degree d. Unsurprisingly, none of the pairs of groups in our experiment were isomorphic and our implementation

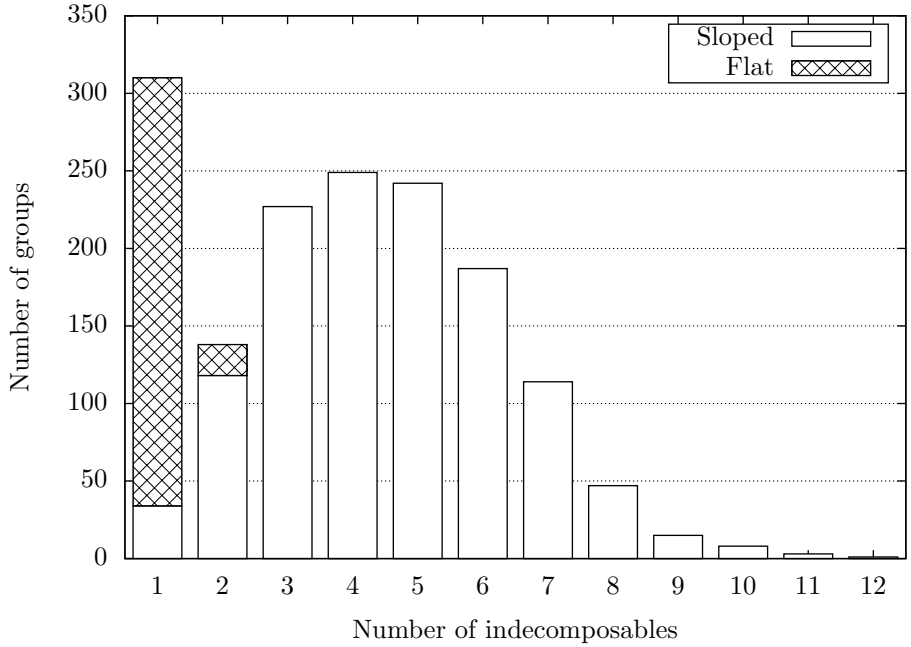

FIGURE 8.2. Bar graph showing the number of indecomposable summands of $A(\circ_G)$ for each group G in Experiment A.

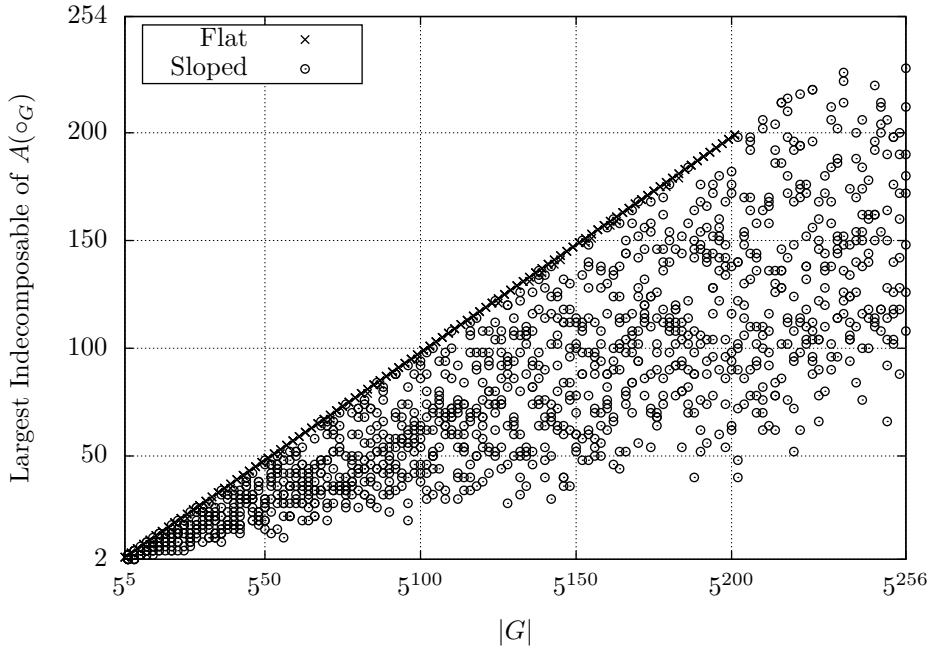

FIGURE 8.3. Graph showing for each group G in Experiment A the dimension of the largest indecomposable summand for $A(\circ_G)$.

quickly determined this by finding that the adjoint algebras $A(\circ_G)$ and $A(\circ_H)$ were non-conjugate.

Experiment C. The lesson we learn from Experiment B is that non-isomorphic groups of genus 2 are usually distinguished by our methods at an early stage. To probe deeper into our algorithm, we conducted a further experiment to produce pairs of groups that are probably non-isomorphic, but which are guaranteed to have conjugate adjoint algebras.

In Experiment C, we again started by constructing a sloped p-group, G, of genus 2 from a random pair of alternating forms. Next, we computed $A = A(\circ_G)$, formed the tensor product $\mathbb{Z}_p \otimes_A \mathbb{Z}_p$, and produced further sloped pairs of forms by taking random 2-dimensional projections of this tensor product. If H is the group built from such a pair of forms, then $A \subseteq A(\circ_H)$ and, with high probability, $A = A(\circ_H)$. Thus, we can generate a stream of groups of genus 2 with fixed adjoint algebras, but having sufficient random variations for testing purposes.

For each $d \in \{2m : 3 \leqslant m \leqslant 9\}$ we created three pairs of groups of order 3^{d+2} and genus 2 (21 pairs of groups in total). For each pair we used the adjoint-tensor method to decide if the groups were isomorphic or non-isomorphic. The test reported that 6 of the 21 pairs were isomorphic groups and the remaining 15 pairs were non-isomorphic.

Verification. As in Theorem 1.1, if our implementation is handed isomorphic groups G and H, the expected affirmative output is accompanied by an explicit isomorphism $\varphi\colon G \to H$. Hence, we have a built-in verification procedure in the isomorphic case. If our implementation is given non-isomorphic groups, on the other hand, the output is simply a boolean confirming that they are not isomorphic without indicating why. While we can turn on simple reporting features that give information such as "adjoint algebras are not conjugate," if we are forced deeper into the algorithm, non-isomorphism is often decided for reasons that are difficult to cast in structural terms. We therefore took the opportunity with the 15 non-isomorphic pairs of groups in Experiment C to seek independent verification using existing MAGMA functions.

For a fixed (allegedly non-isomorphic) pair, we computed the sets of centralizers of all non-central elements in each group, and compared the resulting multisets of subgroups. (We remark that for some pairs this verification took over an hour to complete, in contrast to the seconds taken initially by the adjoint-tensor method to decide non-isomorphism.) This process successfully distinguished 6 of the 15 non-isomorphism claims, but not the other 9 cases. Further testing of the 9 remaining cases was impossible: the group orders were beyond the practical limitations of the default methods in MAGMA, which are based largely on the strategies of [ELGO].

We were not surprised by our inability to use existing tools to distinguish these 9 remaining non-isomorphic pairs. In [LW] a similar construction is used to exhibit an exponentially growing family of pairwise non-isomorphic groups having identical character tables, multisets of centralizers, and a host of other typically decisive isomorphism invariants. Indeed, it is for such families of seemingly indistinguishable groups that algorithms such as the ones reported here are especially valuable.

9. Closing Remarks

Counting. Earlier work confirms that the number of groups of genus 2 grows sufficiently that a classification would be unreasonable [B3, V1, V2]. We can use our methods to prove an exponential bound on the number of these groups.

Proposition 9.1. *The number of pairwise non-isoclinic groups of order p^n that have genus 2 over a field is $p^{n/2+\Theta(1)}$.*

Proof. We have seen that a group of genus 2 over a field $\mathbb{F}_q$ is determined, up to isoclinism, by a pair $\{\Phi_1, \Phi_2\}$ of alternating forms on $\mathbb{F}_q^m$. (For a group of order p^n and genus 2, $n = (2m+2)\log_p q$.) Furthermore, these forms can be written uniquely as an orthogonal sum sloped and flat components. Let $m = s + f$ where s is the dimension of the sloped factor.

To estimate the number of possibilities for the flat part we need only consider the dimensions of the flat indecomposable constituents, forming a partition of the total dimension, f. Furthermore, if $\{\Phi_1, \Phi_2\}$ is a flat indecomposable pair of forms on $\mathbb{F}_q^e$, then $e \geqslant 3$ is odd. Hence, the number of flat indecomposables is at most the number of decompositions $f = \sum_i 2m_i + 1$. This is bounded by the number of partitions of $f/2$, and so is not more than $2^{f/2}$.

We now estimate the possibilities for the sloped portion. We may assume

$$\Psi_1 = \begin{bmatrix} 0 & I_{s/2} \\ -I_{s/2} & 0 \end{bmatrix} \text{ and } \Psi_2 = \begin{bmatrix} 0 & J \\ -J^{\mathrm{tr}} & 0 \end{bmatrix},$$

where J is in generalized Jordan normal form. By a classical result of Frobenius and Hall, the number of Jordan forms (also the number of conjugacy classes in $\mathbb{M}_{s/2}(q)$) is $q^{s/2+o(1)}$. By Theorem 3.22 two sloped pairs determine isomorphic groups if, and only if, they are equivalent under the action of $\Gamma\mathrm{L}(2, \mathbb{F}_q)$. Hence, the total number of pairwise nonisomorphic sloped components is between $q^{s/2-4}$ and $q^{s/2}$.

The total number of pairwise nonisomorphic groups of genus 2 and order p^n is maximized when $\mathbb{F}_q = \mathbb{Z}_p$ and $f \in O(1)$, resulting in the stated estimate. □

We turn next to the degrees of permutation representations having groups of genus 2 as a quotient. This confirms that large groups may be handed to our algorithms for Theorem 1.1 and yet there is a completely deterministic polynomial time solution.

Proposition 9.2. *Let G be a p-group of genus 2 over a field $\mathbb{F}_q$ with fully-refined central decomposition $\{G_1, \cdots, G_\ell\}$. Then G has a faithful representation as a quotient of a permutation group of degree*

$$\deg(G) \leqslant \sum_{i=1}^{\ell} \deg(G_i), \quad \deg(G_i) \leqslant \begin{cases} q^{2c_i \deg a_i(x)}, & H(\mathbb{F}_q[x]/(a_i(x)^{c_i})) \twoheadrightarrow G_i; \\ q^{2m+2}, & G_i \cong H_m^{\flat}(q). \end{cases}$$

Furthermore, $|G| = q^{2s} \prod_{i=1}^{\ell} \deg(G_i)$, where s is the number of G_i that are sloped.

Proof. If $H = H(K)$, K a commutative ring, then the stabilizer of $(1, s, t)$ in H is

$$\left\{ \begin{bmatrix} 1 & 0 & fs \\ 0 & 1 & f \\ 0 & 0 & 1 \end{bmatrix} : f \in K \right\}.$$

Hence, the stabilizers of $(1,0,0)$ and $(1,1,0)$ intersect trivially, so the permutation representation of H on $\{(1,s,t) : s,t \in K\}$ is faithful and transitive of degree $2|K|$. It follows, for each $1 \leqslant i \leqslant \ell$, that

$$\deg(H(\mathbb{F}_q[x]/(a_i(x)^{c_i}))) \leqslant q^{2c_i \deg a_i(x)}.$$

A similar estimate holds for flat indecomposables, but the representation is regular. Since central products are quotients of direct products the claim follows. □

Groups of higher genus. As suggested in Section 4, the adjoint-tensor method is designed to work in much greater generality than genus 2. In [LW], for example, a version of the algorithm handles isomorphism of all quotients of Heisenberg groups over fields in time $O((\log |G|)^6)$. What prevents us from saying more is that one cannot predict the complexity of group isomorphism for a class of groups by the adjoint-tensor method without a priori knowledge of the associated adjoint rings. In the case of quotients of Heisenberg groups $H_m(K)$, K a local Artinian ring, the adjoint rings are generically the same as the adjoint ring of $H_m(K)$, which is none other than $M_{2m}(K)$. (This follows from a Galois correspondence explained in [BW2]). So long as this ring is manageable then some variation of our analysis still applies.

Acknowledgments

The authors wish to thank R. Lipyanski and V. Sergeĭčuk for their most helpful guidance on wild and tame classification problems. We also thank the anonymous referees for some very helpful suggestions that have improved the exposition of the paper in substantive ways.

This work was partially supported by a grant from the Simons Foundation (#281435 to Peter Brooksbank).

References

[B1] Reinhold Baer, *Groups with abelian central quotient group*, Trans. Amer. Math. Soc. **44** (1938), no. 3, 357–386. MR1501972

[B2] Simon R. Blackburn, *Groups of prime power order with derived subgroup of prime order*, J. Algebra **219** (1999), no. 2, 625–657. MR1706841 (2000i:20032)

[B3] James Bond, *Lie algebras of genus one and genus two*, Pacific J. Math. **37** (1971), 591–616. MR0308221 (46 #7336)

[BF] Eva Bayer-Fluckiger, *Principe de Hasse faible pour les systèmes de formes quadratique*, J. Reine Angew. Math. **378** (1987), 53–59. MR0895284 (88g:11015)

[BS] László Babai and Endre Szemerédi, *On the complexity of matrix group problems, I*, Proc. 25th IEEE Sympos. Foundations Comp. Sci., 1984, pp. 229–240.

[BDL$^+$] Genrich Belitskii, Andrii R. Dmytryshyn, Ruvim Lipyanski, Vladimir V. Sergeichuk, and Arkady Tsurkov, *Problems of classifying associative or Lie algebras over a field of characteristic not two and finite metabelian groups are wild*, Electron. J. Linear Algebra **18** (2009), 516–529. MR2538621 (2010i:16024)

[BLS] Genrich Belitskii, Ruvim Lipyanski, and Vladimir Sergeichuk, *Problems of classifying associative or Lie algebras and triples of symmetric or skew-symmetric matrices are wild*, Linear Algebra Appl. **407** (2005), 249–262. MR2161930 (2006i:17014)

[BNV] Simon R. Blackburn, Peter M. Neumann, and Geetha Venkataraman, *Enumeration of finite groups*, Cambridge Tracts in Mathematics, vol. 173, Cambridge University Press, Cambridge, 2007. MR2382539 (2009c:20041)

[BCP] W. Bosma, J. Cannon, and C. Playoust, *The Magma algebra system. I. The user language*, J. Symbolic Comput. **24** (1997), no. 3-4, 235–265. MR1484478

[BL] Peter A. Brooksbank and Eugene M. Luks, *Testing isomorphism of modules*, J. Algebra **320** (2008), no. 11, 4020–4029. MR2464805 (2009h:16001)

[BO] Peter A. Brooksbank and E. A. O'Brien, *Constructing the group preserving a system of forms*, Internat. J. Algebra Comput. **18** (2008), no. 2, 227–241. MR2403820 (2009g:20020)

[BW1] Peter A. Brooksbank and James B. Wilson, *The module isomorphism problem reconsidered*, J. Algebra **421** (2015), 541–559. MR3272396

[BW2] ———, *Groups acting on tensor products*, J. Pure Appl. Algebra **218** (2014), no. 3, 405–416. MR3124207

[BW3] ———, *Intersecting two classical groups*, J. Algebra **353** (2012), 286–297. MR2872448

[BW4] ———, *Computing isometry groups of Hermitian maps*, Trans. Amer. Math. Soc. **364** (2012), no. 4, 1975–1996. MR2869196

[CdGS] Serena Cicalò, Willem A. de Graaf, and Csaba Schneider, *Six-dimensional nilpotent Lie algebras*, Linear Algebra Appl. **436** (2012), no. 1, 163–189. MR2859920

[D] Jean Dieudonné, *Sur la réduction canonique des couples de matrices*, Bull. Soc. Math. France **74** (1946), 130–146. MR0022826 (9,264f)

[ELGO] Bettina Eick, C. R. Leedham-Green, and E. A. O'Brien, *Constructing automorphism groups of p-groups*, Comm. Algebra **30** (2002), no. 5, 2271–2295. MR1904637 (2003d:20027)

[F] A. A. Finogenov, *On finite p-groups with a cyclic commutator group and cyclic center*, Mat. Zametki **63** (1998), no. 6, 911–922; English transl., Math. Notes **63** (1998), no. 5-6, 802–812. MR1679224 (2000a:20044)

[GG] Daniel Goldstein and Robert M. Guralnick, *Alternating forms and self-adjoint operators*, J. Algebra **308** (2007), no. 1, 330–349. MR2290925 (2008b:20050)

[H] P. Hall, *The classification of prime-power groups*, J. Reine Angew. Math. **182** (1940), 130–141. MR0003389

[HEO] Derek F. Holt, Bettina Eick, and Eamonn A. O'Brien, *Handbook of computational group theory*, Discrete Mathematics and its Applications, Chapman & Hall/CRC, Boca Raton, 2005. MR2129747 (2006f:20001)

[K] M. S. Knebelman, *Classification of Lie algebras*, Ann. of Math. (2) **36** (1935), no. 1, 46–56. MR1503207

[KL] W. M. Kantor and E. M. Luks, *Computing in Quotient Groups*, Proceedings of the Twenty-second Annual ACM Symposium on Theory of Computing, 1990, pp. 524–534.

[L1] Y. K. Leong, *Odd order nilpotent groups of class two with cyclic centre*, J. Austral. Math. Soc. **17** (1974), 142–153. Collection of articles dedicated to the memory of Hanna Neumann, VI. MR0347972 (50 #470)

[L2] E. M. Luks, *Computing in solvable matrix groups*, 33rd Annual Symposium on Foundations of Computer Science, Pittsburgh, Oct. 24–27 (1992), 111–120.

[LGM] C. R. Leedham-Green and S. McKay, *The structure of groups of prime power order*, London Mathematical Society Monographs. New Series, vol. 27, Oxford University Press, Oxford, 2002. Oxford Science Publications. MR1918951 (2003f:20028)

[LGS1] C. R. Leedham-Green and Leonard H. Soicher, *Collection from the left and other strategies*, J. Symbolic Comput. **9** (1990), 5–6,665–675. MR1075430 (92b:20021)

[LGS2] ———, *Symbolic collection using Deep Thought*, LMS J. Comput. Math. **1** (1998), 9–24. MR1635719 (99f:20002)

[LW] Mark L. Lewis and James B. Wilson, *Isomorphism in expanding families of indistinguishable groups*, Groups Complex. Cryptol. **4** (2012), no. 1, 73–110. MR2921156

[LM] E.M. Luks and T. Miyazaki, *Polynomial-time normalizers*, Discrete Math. Theor. Comput. Sci. **4** (2011), 61–96. MR2862561 (20010)

[M1] Gary L. Miller, *Graph isomorphism, general remarks*, Journal of Computer and System Sciences **18** (1979), no. 2, 128 - 142. MR0532172 (80h:68056)

[M2] V. V. Morozov, *Classification of nilpotent Lie algebras of sixth order*, Izv. Vysš. Učebn. Zaved. Matematika **1958** (1958), no. 4 (5), 161–171. MR0130326 (24 #A190)

[MH] John Milnor and Dale Husemoller, *Symmetric bilinear forms*, Springer-Verlag, New York-Heidelberg, 1973. Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 73. MR0506372 (58 #22129)

[N] Peter M. Neumann, *Some algorithms for computing with finite permutation groups*, Proceedings of groups—St. Andrews 1985, London Math. Soc. Lecture Note Ser., vol. 121, Cambridge Univ. Press, Cambridge, 1986, pp. 59–92. MR896501

[O] E. A. O'Brien, *Isomorphism testing for p-groups*, J. Symbolic Comput. **17** (1994), no. 2, 131, 133–147. MR1283739 (95f:20040b)

[P] Donald S. Passman, *A course in ring theory*, The Wadsworth & Brooks/Cole Mathematics Series, Wadsworth & Brooks/Cole Advanced Books & Software, Pacific Grove, CA, 1991. MR1096302

[RW] D. Rosenbaum and F. Wagner, *Beating the Generator-Enumeration Bound for p-Group Isomorphism*, Theoretical Computer Science **593** (2015), 16-25.

[S1] Rudolf Scharlau, *Paare alternierender Formen*, Math. Z. **147** (1976), no. 1, 13–19. MR0419484 (54 #7505)

[S2] Ruth Schwingel, *Two matrix group algorithms with applications to computing the automorphism group of a finite p-group*, PhD thesis, Queen Mary, University of London, 2000.

[S3] Ákos Seress, *Permutation group algorithms*, Cambridge Tracts in Mathematics, vol. 152, Cambridge University Press, Cambridge, 2003. MR1970241 (2004c:20008)

[S4] Arne Storjohann, *An $O(n^3)$ algorithm for the Frobenius normal form*, Proceedings of the 1998 International Symposium on Symbolic and Algebraic Computation (Rostock), ACM, New York, 1998, pp. 101–104 (electronic). MR1805172

[vzGG] Joachim von zur Gathen and Jürgen Gerhard, *Modern computer algebra*, 2nd ed., Cambridge University Press, Cambridge, 2003. MR2001757 (2004g:68202)

[V1] A. L. Vishnevetskiĭ, *Groups of class 2 and exponent p with commutant of order p^2*, Dokl. Akad. Nauk Ukrain. SSR Ser. A **9** (1980), 9–11, 103. MR593560 (82d:20026)

[V2] ———, *A system of invariants of certain groups of class 2 with commutator subgroup of rank two*, Ukrain. Mat. Zh. **37** (1985), no. 3, 294–300, 403. MR795568 (86k:20033)

[W1] James B. Wilson, *Decomposing p-groups via Jordan algebras*, J. Algebra **322** (2009), no. 8, 2642–2679. MR2559855 (2010i:20016)

[W2] ———, *Division, adjoints, and dualities of bilinear maps*, Comm. Algebra **41** (2013), no. 11, 3989–4008. MR3169502

[W3] ———, *Optimal algorithms of Gram-Schmidt type*, Linear Algebra Appl. **438** (2013), no. 12, 4573–4583. MR3039211

[W4] ———, *Existence, algorithms, and asymptotics of direct product decompositions, I*, Groups Complex. Cryptol. **4** (2012), no. 1, 33–72. MR2921155

[W5] ———, *Finding central decompositions of p-groups*, J. Group Theory **12** (2009), no. 6, 813–830. MR2582050 (2011a:20044)

Department of Mathematics, Bucknell University, Lewisburg, PA 17837, USA
E-mail address: pbrooksb@bucknell.edu

Department of Mathematics, Colorado State University, Fort Collins, CO 80523, USA
E-mail address: maglione@math.colostate.edu

Department of Mathematics, Colorado State University, Fort Collins, CO 80523, USA
E-mail address: James.Wilson@ColoState.Edu