



On p -Group Isomorphism: search-to-decision, counting-to-decision, and nilpotency class reductions via tensors

JOSHUA A. GROCHOW, Departments of Computer Science and Mathematics, University of Colorado Boulder, USA

YOUMING QIAO, Centre for Quantum Software and Information, University of Technology Sydney, Australia

In this paper we study some classical complexity-theoretic questions regarding GROUP ISOMORPHISM (GPI). We focus on p -groups (groups of prime power order) with odd p , which are believed to be a bottleneck case for GPI, and work in the model of matrix groups over finite fields. Our main results are as follows.

- Although search-to-decision and counting-to-decision reductions have been known for over four decades for GRAPH ISOMORPHISM (GI), they had remained open for GPI, explicitly asked by Arvind & Torán (*Bull. EATCS*, 2005). Extending methods from TENSOR ISOMORPHISM (Grochow & Qiao, ITCS 2021), we show moderately exponential-time such reductions within p -groups of class 2 and exponent p .
- Despite the widely held belief that p -groups of class 2 and exponent p are the hardest cases of GPI, there was no reduction to these groups from *any* larger class of groups. Again using methods from TENSOR ISOMORPHISM (*ibid.*), we show the first such reduction, namely from isomorphism testing of p -groups of “small” class and exponent p to those of class two and exponent p .

For the first results, our main innovation is to develop linear-algebraic analogues of classical graph coloring gadgets, a key technique in studying the structural complexity of GI. Unlike the graph coloring gadgets, which support restricting to various subgroups of the symmetric group, the problems we study require restricting to various subgroups of the general linear group, which entails significantly different and more complicated gadgets. The analysis of one of our gadgets relies on a classical result from group theory regarding random generation of classical groups (Kantor & Lubotzky, *Geom. Dedicata*, 1990). For the nilpotency class reduction, we combine a runtime analysis of the Lazard correspondence with TENSOR ISOMORPHISM-completeness results (Grochow & Qiao, *ibid.*).

CCS Concepts: • **Theory of computation** → **Problems, reductions and completeness**.

Additional Key Words and Phrases: group isomorphism, search-to-decision, counting-to-decision, p -groups

1 INTRODUCTION

In this paper, we study the algorithmic problem of deciding whether two finite groups are isomorphic, known as the GROUP ISOMORPHISM problem (GPI). Different variants of the GPI problem arise, with correspondingly different complexities, when the groups are given in different ways, e.g. by a generating set of permutations, a generating set of matrices, a full multiplication table, or a black box oracle. In its various incarnations, GPI is a fundamental problem in computational algebra and computational complexity. The generator-enumerator

Authors' addresses: Joshua A. Grochow, Departments of Computer Science and Mathematics, University of Colorado Boulder, Boulder, Colorado, USA, 80309-0430, jgrochow@colorado.edu; Youming Qiao, Centre for Quantum Software and Information, University of Technology Sydney, Sydney, New South Wales, Australia, 2007, jimmyqiao86@gmail.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1942-3454/2023/9-ART \$15.00

<https://doi.org/10.1145/3625308>

algorithm solves isomorphism in $|G|^{\log |G| + O(1)}$ -time [29, 60]¹, and even the current state of the art for general groups—in any of the aforementioned input models—is still $|G|^{\Theta(\log |G|)}$ [10, 11, 18, 28, 52, 68, 72]. Nonetheless, over the past 15 years there has been significant progress on efficient isomorphism tests in various classes of groups: here is an incomplete list of references [5–7, 13, 14, 16, 33, 34, 50, 51, 65, 67, 68].

When given by multiplication tables, GPI reduces to GI [75], and in the other, more realistic (for computer algebra systems) and more succinct models, we get a reduction in the other direction [35, 37, 54, 59]. As a result, the techniques and complexity of GPI are closely bound up with GI. However, since the techniques used in GPI are often independent of the input model, we are free to focus on the abstract structure of the groups in question, and the choice of input model is then essentially just a choice of how we measure and report the running time. For example, if GI is in P, then GPI can be solved in $\text{poly}(|G|)$ time [75]; if GPI for groups given by a generating set of m matrices of size $n \times n$ over $\mathbb{F}_p$ can be solved in $p^{O(n+m)}$ time, then GI is in P [59] (see [37] for a simplified reduction).

For GI, a wide variety of algorithmic and structural complexity results are known (see, e.g., [4, 36, 47]). In particular, there are polynomial-time search-to-decision and counting-to-decision reductions [56], so search, counting, and decision are all equivalent for GI. (This was an early piece of evidence that GI was not likely to be NP-complete, since for NP-complete problems, their counting variants are typically #P-complete, hence at least as hard as all of PH [70].) For GPI, no such reductions are known, even in restricted classes of groups; Arvind and Torán [3, Problem 16] explicitly asked for such reductions. Additionally, for GI, there are many classes of graphs for which the isomorphism problem remains GI-complete—such as graphs of diameter 2 and radius 1, directed acyclic graphs, regular graphs, line graphs, polytopal graphs [75]—but no such analogous results are known for GPI.

In this paper, we make progress on all three of these questions, within the class of groups widely believed to be hardest cases of GPI, namely the p -groups of nilpotency class 2 and exponent p ; these are groups of order a power of the prime p , such that G modulo its center is abelian, and such that $g^p = 1$ for all $g \in G$. (Throughout most of this paper we assume p is an odd prime.) For each of our three main results, we now give further motivation before stating it formally.

1.1 Main results

Search-to-decision reductions. The “decision versus search” question is a classical one in complexity theory, having attracted the attention of researchers since the introduction of NP. Efficient search-to-decision reductions for SAT and GI are now standard. Valiant first showed the existence of an NP *relation* for which search does not reduce to decision in polynomial time [71]. A celebrated result of Bellare and Goldwasser shows that, assuming $\text{DTIME}(2^{2^{O(n)}}) \neq \text{NTIME}(2^{2^{O(n)}})$, there exists an NP *language* for which search does not reduce to decision in polynomial time [9]. However, as usual for such statements based on complexity-theoretic assumptions, the problems constructed by such a proof are considered somewhat unnatural, and natural problems for which search seems not reducible to decision are rare. The most famous candidate may be FACTORING (with the decision version being PRIMALITY)² and NASH EQUILIBRIUM [19] (the decision version is trivial).

Theorem A. *Let p be an odd prime, and let $\text{GPIso2Exp}(p)$ denote the isomorphism problem for p -groups of class 2 and exponent p in the model of matrix groups over $\mathbb{F}_p$. For groups of order p^n , there is a search-to-decision reduction for $\text{GPIso2Exp}(p)$ running in time $p^{O(n)} = \text{poly}(|G|)$.*

¹Miller [60] attributes this algorithm to Tarjan.

²Here we are thinking of FACTORING as the search problem corresponding to the relation $\{(n, d) : d \text{ is a proper divisor of } n\} \subseteq \mathbb{N} \times \mathbb{N}$, so that the existence problem is then precisely PRIMALITY.

We note that this improves over the “brute-force” generator-enumerator technique, which runs in time $p^{\Theta(n^2)} = |G|^{\Theta(\log |G|)}$.

Remark 1.1. Nearly all our results about groups require p to be an *odd* prime (many of our results on tensors or matrix spaces should still work when $p = 2$). There are (at least) two crucial differences in the $p = 2$ case for groups. The first is that for 2-groups, the Baer correspondence no longer works in the form presented here (rather, there is a different correspondence involving 2-cocycles and quadratic maps rather than bilinear maps). The second issue is that groups of exponent 2 are all Abelian; the smallest-exponent non-Abelian 2-groups are of exponent 4. If one then translates between groups and tensors, one would get tensors over the ring $\mathbb{Z}/4\mathbb{Z}$. As $\mathbb{Z}/4\mathbb{Z}$ is no longer a field, compared to our setting where we get to work over $\mathbb{Z}/p\mathbb{Z}$, this introduces significant additional complications. We leave working with such groups and tensors to future work.

We note that $\text{GPIso2Exp}(p)$ seems different from all the problems listed above in terms of search-to-decision reductions, in the following ways. First, unlike SAT (propositional Boolean satisfiability) and GI, a polynomial-time search-to-decision reduction has been open for decades, whereas those for SAT and GI are straightforward. Note that a polynomial-time reduction would need to run in time $\text{poly}(n, \log p)$, and we find it unlikely that the time complexity of our reduction can be brought down this far with current techniques. Second, unlike FACTORING and NASH EQUILIBRIUM, whose decision versions are computationally easy (PRIMALITY is easily seen to be in $\text{RP} \cap \text{coNP}$, even if the proof it is in P [1] is quite nontrivial, and the decision version of NASH EQUILIBRIUM has a trivial “yes” answer by Nash’s Theorem), its decision version also seems to require deeper techniques. Indeed, it is a long-standing open problem to test isomorphism of p -groups of class 2 and exponent p in time polynomial in the group order, which already can be exponential in the input size if the input is given by a generating set of matrices.

Counting-to-decision reductions. Counting-to-decision reductions are also of great interest in complexity theory. An efficient counting-to-decision reduction for GI is also a well-known result [56]. In contrast, for SAT, a polynomial-time counting-to-decision reduction would imply that PH collapses [70].

Theorem B. *For p an odd prime, $p \geq n^{\Omega(1)}$, there is a randomized counting-to-decision reduction for $\text{GPIso2Exp}(p)$ for groups of order p^n , running in time $p^{O(n)} = \text{poly}(|G|)$.*

As with Theorem A, this improves the previous-best “brute-force” $p^{O(n^2)} = |G|^{O(\log |G|)}$.

Also as in the case of search-to-decision, $\text{GPIso2Exp}(p)$ seems different from the problems listed above in terms of reducing counting to decision. First, a polynomial-time counting-to-decision reduction for $\text{GPIso2Exp}(p)$ remains open after 40 years of studying GPI (going back at least to [29, 60]), whereas the reduction for GI was found within the first decade of the rise of computational complexity theory. Second, unlike SAT, for which there have been no non-trivial algorithms to reduce exact counting to decision, we show a moderately exponential-time algorithm for $\text{GPIso2Exp}(p)$. As Ryan Williams pointed out to us, asking for the existence of a subexponential-time counting-to-decision reduction for SAT seems to lead to asking for the relation between the decision [38] and the counting [25] versions of the Exponential Time Hypothesis.

Nilpotency class reduction. Unlike the case of GRAPH ISOMORPHISM, for GPI essentially the only class of groups for which isomorphism is known to be as hard as the general case are those which are directly indecomposable, that is, they cannot be written as a direct product $A \times B$ with both A, B nontrivial [45, 73, 74]. However, this result is the group analogue of saying that isomorphism of connected graphs is GI-complete, so although useful (and much less trivial than in the case of graphs vs connected graphs), from a structural perspective it is more like a zero-th step.

For a variety of reasons (e.g., [32]), p -groups of nilpotency class 2 and exponent p are widely believed to be the hardest cases of GPI, but to date there is no known reduction from isomorphism in *any* larger class of groups to

this class. The TENSOR ISOMORPHISM-completeness of testing isomorphism in this class of groups (when given by generating matrices over $\mathbb{F}_p$) suggests an additional reason for hardness [35] (see also Section 6.1). Here, we leverage that completeness result to give a reduction within GrP itself. While it falls short of being GrP-complete (equivalent to GrP), this is the first such reduction that we are aware of.

To state our result, we need to first recall the definition of nilpotency class. We will give an inductive definition: a group G is nilpotent of class 1 if it is abelian, and nilpotent of class $c > 1$ if $G/Z(G)$ (G modulo its center) is nilpotent of class $c - 1$. Recall that a finite group is nilpotent iff it is the direct product of its Sylow p -subgroups, so from the comment above, isomorphism of nilpotent groups is polynomial-time equivalent to isomorphism of p -groups (for varying p).

Theorem P. *Let p be an odd prime. For groups given by generating sets of m matrices of size $n \times n$ over $\mathbb{F}_{p^e}$, GROUP ISOMORPHISM for p -groups of exponent p and class $c < p$ reduces to GROUP ISOMORPHISM for p -groups of exponent p and class 2 in time $\text{poly}(n, m, e \log p)$.*

In fact, because the Lazard correspondence works whenever all subgroups generated by 3 elements have nilpotency class $< p$, our reduction also works in this more general setting. For example, as a consequence of Thm. P, testing isomorphism of 5-groups in which every 3-generated subgroup has class 4 (the groups themselves may have larger class) reduces to testing isomorphism of 5-groups of class 2 in the matrix group model over fields of characteristic 5.

Remark 1.2. Two additional results would suffice to get the analogous result in the Cayley table model. The first is to compute the Lazard correspondence in the Cayley table model in time $\text{poly}(|G|)$; we thank an anonymous ITCS reviewer for pointing out that this can be achieved by applying the matrix Lazard correspondence (see Proposition 6.4) to the left regular representation of the group on itself. The second is to improve the blow-up in the reduction from (LIE) ALGEBRA ISOMORPHISM to 3TI from [31]. Currently this reduction increases the dimension quadratically, which means the size of the group becomes $|G|^{O(\log |G|)}$ after the reduction; instead, we would need a reduction that increases the dimension only linearly.

Remark 1.3. One may also ask whether our theorems can be combined, in order to get search-to-decision and counting-to-decision reductions for p -groups of class $c < p$ instead of only class 2. We believe this should be approachable, but again the quadratic increase in dimension in reductions, mentioned in the previous remark, gets in the way. The quadratic increase makes the square-root exponential reductions into ordinary exponential reductions, negating any gains.

1.2 Main techniques and proof strategies

All our results are based on the connection with TENSOR ISOMORPHISM (TI) [35]. Let $\Lambda(n, \mathbb{F})$ denote the space of $n \times n$ skew-symmetric (alternating) matrices over $\mathbb{F}$. Then the Baer Correspondence [8] gives an equivalence between

$$\left\{ \begin{array}{l} p\text{-groups of class 2, exponent } p, \\ G/Z(G) \cong \mathbb{Z}_p^n, Z(G) \cong \mathbb{Z}_p^m \end{array} \right\} \longleftrightarrow \left\{ \begin{array}{l} \mathcal{A} \leq \Lambda(n, \mathbb{F}_p) \\ \dim \mathcal{A} = m \end{array} \right\} \longleftrightarrow \left\{ \begin{array}{l} \text{Nilpotent } \mathbb{F}_p\text{-Lie algebras of} \\ \text{class 2, } L/Z(L) \cong \mathbb{F}_p^n, Z(L) \cong \mathbb{F}_p^m \end{array} \right\}$$

in such a way that two such groups are isomorphic iff the corresponding Lie algebras are isomorphic iff the corresponding matrix spaces $\mathcal{A}, \mathcal{B} \leq \Lambda(n, \mathbb{F}_p)$ are isometric. Here, we say that two such linear subspaces are *isometric* if there is an invertible matrix $L \in \text{GL}(n, \mathbb{F}_p)$ such that $\mathcal{B} = L^t \mathcal{A} L := \{L^t A L : A \in \mathcal{A}\}$. The corresponding computational problem is:

Definition 1.4 (The ALTERNATING MATRIX SPACE ISOMETRY problem).

Input: $A_1, \dots, A_m$ and $B_1, \dots, B_m$, $n \times n$ alternating³ matrices over a field $\mathbb{F}$,

Decide: Is there a $L \in \text{GL}(n, \mathbb{F})$, such that the linear span of $\{A_i : i \in [m]\}$ is equal to the linear span of $\{L^t B_i L : i \in [m]\}$?

Our search- and counting-to-decision reductions (Thms. A and B) actually follow from analogous results on ALTERNATING MATRIX SPACE ISOMETRY (Thms. A' and B'), using a constructive version of the Baer Correspondence communicated to us by James B. Wilson (Lem. 6.2). The viewpoint of alternating matrix spaces made the constructions much easier to find and reason about.

Our nilpotency class reduction uses a constructive version of the Lazard Correspondence (Prop. 6.4), which generalizes the Baer correspondence to nilpotency class $c < p$; the TI-completeness of LIE ALGEBRA ISOMORPHISM for nilpotent Lie algebras of class 2 (a combination of reductions from [31] and [35]); and finally the aforementioned constructive Baer Correspondence to go back to p -groups of class 2.

In the remainder of this section we give more details of the techniques involved.

1.2.1 Linear algebraic coloring gadgets. Our most novel technique is to devise linear algebraic analogues for ALTERNATING MATRIX SPACE ISOMETRY of the graph coloring gadget, a key technique in the structural complexity study of GRAPH ISOMORPHISM (see, e. g., [47]). This technique is crucial in the following theorems, used to prove Thms. A and B, respectively.

Theorem A'. *Let q be a prime power. There is a search-to-decision reduction for ALTERNATING MATRIX SPACE ISOMETRY which, given $n \times n$ alternating matrix spaces $\mathcal{A}, \mathcal{B}$ over $\mathbb{F}_q$ of dimension m , computes an isometry between them if they are isometric, in time $q^{\tilde{O}(n)}$ or in time $q^{O(n+m)}$. The reduction queries the decision oracle with inputs of dimension at most $O(n^2)$.*

Theorem B'. *For q a prime power with $q = n^{\Omega(1)}$, there is a randomized counting-to-decision reduction for ALTERNATING MATRIX SPACE ISOMETRY which, given $n \times n$ alternating matrix spaces $\mathcal{A}, \mathcal{B}$ over $\mathbb{F}_q$ of dimension m , computes the number of isometries from $\mathcal{A}$ to $\mathcal{B}$ in time $q^{O(n)}$. The reduction queries the decision oracle with inputs of dimension at most $O(n^2)$.*

Let us first briefly review the graph coloring gadgets. Suppose we have a graph $G = (V, E)$ with the vertices colored, i. e., there is a map $f : V \rightarrow \{1, \dots, c\} =: [c]$, where we view $[c]$ as the set of colors. Let $n = |V|$. Suppose we want to construct an uncolored graph $\tilde{G}$, in which the color information carried by f is encoded. One way to achieve this is the following. (See [47] for other more efficient constructions.) For every $v \in V$, if $v \in V$ is assigned color $k \in [c]$, then attach a “star” of size kn to v , that is add kn new vertices to G and attach them all to v . We then get a graph $\tilde{G}$ with $O(cn^2)$ vertices, and we see that an automorphism of $\tilde{G}$, when restricting to V , has to map $v \in V$ to another $v' \in V$ of the same color, as degrees need to be preserved under automorphisms.

Such an idea can be carried out in the 3-tensor context as in [31], but with a significant loss of efficiency, which prevents its use for search- and counting-to-decision reductions and indicates the needs for new techniques. To illustrate the situation, we consider a toy problem. To ease the presentation, we adopt a perspective on 3-tensors that we hope is clear on its own; the analogy with the graph case is fairly close, but not immediately obvious, and we present it in full detail in Sec. 3. Note that by slicing a 3-tensor along one direction, we get a tuple of matrices (see also Section 2); in the following of this subsection we shall mostly work with matrix tuples.

Let $\mathbf{A} = (A_1, \dots, A_m) \in \text{M}(n, \mathbb{F})^m$ be a tuple of matrices, where A_i 's are linearly independent, and $\text{M}(n, \mathbb{F})$ denotes the space of $n \times n$ matrices over $\mathbb{F}$. There are two natural actions on $\mathbf{A}$. The first action is $S = (s_{i,j}) \in \text{GL}(m, \mathbb{F})$ on $\mathbf{A}$ by sending A_j to $\sum_{i \in [m]} s_{i,j} A_i$. Denote the resulting matrix tuple by $\mathbf{A}^S$. The second action is $(L, R) \in \text{GL}(n, \mathbb{F}) \times \text{GL}(n, \mathbb{F})$ on $\mathbf{A}$ by sending A_j to $LA_j R^t$ for $j = 1, \dots, m$. Denote the resulting matrix tuple

³An $n \times n$ matrix A over $\mathbb{F}$ is alternating if for every $v \in \mathbb{F}^n$, $v^t A v = 0$. When $\mathbb{F}$ is not of characteristic 2, this is equivalent to being skew-symmetric $A^t = -A$.

by LAR^t . For two tuples $\mathbf{A}, \mathbf{B}$, and for the purposes of this illustration, let us define the set of isomorphisms as $\text{Iso}(\mathbf{A}, \mathbf{B}) = \{S \in \text{GL}(m, \mathbb{F}) : \exists L, R \in \text{GL}(n, \mathbb{F}), LAR^t = \mathbf{B}^S\}$.

In the counting-to-decision reduction we will need to test isomorphism of such tuples under the action by *diagonal* matrices. Let $\text{diag}(m, \mathbb{F})$ denote the subgroup of $\text{GL}(m, \mathbb{F})$ consisting of diagonal matrices. Our goal then is to construct $\tilde{\mathbf{A}} = (\tilde{A}_1, \tilde{A}_2, \tilde{A}_3) \in M(N, \mathbb{F})^3$ and $\tilde{\mathbf{B}}$, such that $\text{Iso}(\tilde{\mathbf{A}}, \tilde{\mathbf{B}}) = \text{Iso}(\mathbf{A}, \mathbf{B}) \cap \text{diag}(3, \mathbb{F})$. The construction we use, from [31], is as follows. Let $N = 2^3 \cdot n = 8n$, and let

$$\tilde{A}_1 = \begin{bmatrix} A_1 & 0 & 0 & 0 \\ 0 & I_n & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \tilde{A}_2 = \begin{bmatrix} A_2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & I_{2n} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \tilde{A}_3 = \begin{bmatrix} A_3 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & I_{4n} \end{bmatrix}, \quad (1)$$

where I_s denotes the identity matrix of size s , and 0's denote all-zero matrices of appropriate sizes, and define $\tilde{\mathbf{B}}$ similarly. By [31, Lemma 2.2], we have $\text{Iso}(\tilde{\mathbf{A}}, \tilde{\mathbf{B}}) = \text{Iso}(\mathbf{A}, \mathbf{B}) \cap \text{diag}(3, \mathbb{F})$. The proof, while not difficult, relies on certain algebraic machineries like the Krull–Schmidt Theorem for quiver representations. For our purpose, we only point out that a key in the proof is that $\text{Iso}(\tilde{\mathbf{A}}, \tilde{\mathbf{B}}) \subseteq \text{diag}(3, \mathbb{F})$, which can be easily checked by comparing the ranks of the $\tilde{A}_i, \tilde{B}_i$.

The preceding gadget construction can be generalized to handle subgroups of $\text{GL}(n, \mathbb{F})$ of the form $\left\{ \begin{bmatrix} S_1 & 0 & \dots & 0 \\ 0 & S_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & S_c \end{bmatrix} : S_i \in \text{GL}(n_i, \mathbb{F}) \right\}$, where $c = O(\log n)$. We shall refer to this gadget as the Futorny–Grochow–Sergeichuk gadget, or FGS gadget for short.

However, the FGS gadget cannot be used for search- and counting-to-decision reductions in Thms. A and B. The key bottleneck is the restriction that $c = O(\log n)$. To check why this is so reveals an interesting distinction between the combinatorial and the linear algebraic worlds. Recall that in the graph setting, if there are c colors, we need stars of size at most cn . While in the linear algebraic setting, if there are c components, the biggest identity matrix needs to be of size $2^c \cdot n \times 2^c \cdot n$. The reason is that we can do non-trivial linear combinations of the matrices $\tilde{A}_i$, so several matrices of small ranks might be combined to get a matrix of large rank. Indeed, in Eq. 1, if $\tilde{A}_3$ was accompanied with I_{3n} instead of I_{4n} , then a non-trivial linear combination of $\tilde{A}_1$ and $\tilde{A}_2$ could be of rank the same as $\tilde{A}_3$, and the argument that $\text{Iso}(\tilde{\mathbf{A}}, \tilde{\mathbf{B}}) \subseteq \text{diag}(m, \mathbb{F})$ would not go through. That's why we need such exponential growth as the number of components grow.

To address this challenge, we devise two new gadgets, which restrict to the monomial group and the diagonal group, respectively.

The monomial group of $\text{GL}(n, \mathbb{F})$, denoted as $\text{Mon}(n, \mathbb{F})$, consists of monomial matrices, i.e. a matrix with exactly one non-zero entry in each row and each column. We design a gadget that restricts to $\text{Mon}(n, \mathbb{F})$, which is the key in the search-to-decision reduction (Theorem A').

In the case of $\mathbb{F} = \mathbb{F}_q$ and $q = n^{\Omega(1)}$, we design a gadget that restricts to $\text{diag}(n, q)$, which is the key in the counting-to-decision reduction (Theorem B'). The gadget for restricting to monomial groups cannot be used in the counting-to-decision reduction. Its construction is already delicate, and the analysis is involved, relying on a celebrated result of Kantor and Lubotzky regarding random generation of classical groups [44].

1.2.2 Constructive Lazard correspondence. In light of the TI-completeness of isomorphism of class 2 p -groups given by matrices over finite fields of characteristic p [35], the key idea here is how to reduce isomorphism for other classes of groups to some tensor problem. For groups in general this seems quite difficult, as tensors are multilinear and groups are fundamentally not. But for p -groups of nilpotency class $< p$, the Lazard correspondence gives an equivalence between the category of such groups and a corresponding category of Lie algebras (over the same field, nilpotent of the same class). If we could make this correspondence computationally efficient, we

would then be in the fortunate setting in which LIE ALGEBRA ISOMORPHISM is multilinear, and is in TI [31], so we can then reduce back to isomorphism of class 2 p -groups. We observe (Proposition 6.4) that when the groups are given by matrices in characteristic p , the Lazard correspondence can be efficiently computed using the usual matrix logarithm and exponential.

The restriction to groups of nilpotency class $c < p$ comes entirely from the Lazard correspondence, which is also known only to work under this same assumption (see [62] for details, and what can be said when $c = p$, but unfortunately already when $c = p$ one no longer gets an equivalence up to isomorphism). Despite this restriction, we note that we know of no prior reductions from *any* class of groups to p -groups of class 2.

In Rmk. 1.2 we discuss the ingredients necessary to get the same result for GrI in the Cayley table model, which seems approachable.

1.3 Organization of the paper

In Section 2 we present preliminaries and notation. In Section 3 we present more details of the analogy with individualizing vertices in graphs by attaching stars, using the example of reducing MONOMIAL CODE EQUIVALENCE to TENSOR ISOMORPHISM. In Section 4 we present our gadget to restrict to the monomial subgroup, an example use of this to reduce GI to ALTERNATING MATRIX SPACE ISOMETRY, and Thm. A'. In Section 5 we prove Thm. B'. In Section 6 we present the constructive Baer and Lazard Correspondences, and use them to derive Thms. A and B from Thms. A' and B', respectively, as well as proving Thm. P. Finally, in Section 7 we conclude with open questions and discuss the relationship between this work and the authors' line of work on TENSOR ISOMORPHISM.

2 PRELIMINARIES

Font	Object	Space of objects
$A, B, \dots$	matrix	$M(n, \mathbb{F})$ or $M(\ell \times n, \mathbb{F})$
$\mathbf{A}, \mathbf{B}, \dots$	matrix tuple	$M(n, \mathbb{F})^m$ or $M(\ell \times n, \mathbb{F})^m$
$\mathcal{A}, \mathcal{B}, \dots$	matrix space	[Subspaces of $M(n, \mathbb{F})$ or $\Lambda(n, \mathbb{F})$]
$A, B, \dots$	3-way array	$T(\ell \times n \times m, \mathbb{F})$

Table 1. Summary of notation related to 3-way arrays and tensors.

Vector spaces. Let $\mathbb{F}$ be a field. In this paper we only consider finite-dimensional vector spaces over $\mathbb{F}$. We use $\mathbb{F}^n$ to denote the vector space of length- n *column* vectors. The i th standard basis vector of $\mathbb{F}^n$ is denoted $\vec{e}_i$. Depending on the context, $\mathbf{0}$ may denote the zero vector space, a zero vector, or an all-zero matrix. For S a set of vectors, we use $\langle S \rangle$ to denote the subspace spanned by elements in S .

Some groups. The general linear group of degree n over a field $\mathbb{F}$ is denoted by $GL(n, \mathbb{F})$. The symmetric group of degree n is denoted by S_n . The natural embedding of S_n into $GL(n, \mathbb{F})$ is to represent permutations by permutation matrices. The subgroup of $GL(n, \mathbb{F})$ consisting of diagonal matrices is called the *diagonal subgroup*, denoted by $\text{diag}(n, \mathbb{F})$. A *monomial matrix* is a product of a diagonal and a permutation matrix; equivalently, each row and each column has exactly one non-zero entry. The collection of monomial matrices forms a subgroup of $GL(n, \mathbb{F})$, which we call the *monomial subgroup* and denote by $\text{Mon}(n, \mathbb{F})$. It is the semi-direct product $\text{diag}(n, \mathbb{F}) \rtimes S_n \cong (\mathbb{F}^*)^n \rtimes S_n$.

Nilpotent groups. If A, B are two subsets of a group G , then $[A, B]$ denotes the subgroup generated by all elements of the form $[a, b] = aba^{-1}b^{-1}$, for $a \in A, b \in B$. The *lower central series* of a group G is defined as follows: $\gamma_1(G) = G, \gamma_{k+1}(G) = [\gamma_k(G), G]$. A group is *nilpotent* if there is some c such that $\gamma_{c+1}(G) = 1$; the smallest such c is called the *nilpotency class* of G , or sometimes just “class” when it is understood from context. A finite group

is nilpotent if and only if it is the product of its Sylow subgroups; in particular, all groups of prime power order are nilpotent.

Matrices. Let $M(\ell \times n, \mathbb{F})$ be the linear space of $\ell \times n$ matrices over $\mathbb{F}$, and $M(n, \mathbb{F}) := M(n \times n, \mathbb{F})$. Given $A \in M(\ell \times n, \mathbb{F})$, A^t denotes the transpose of A .

A matrix $A \in M(n, \mathbb{F})$ is *alternating*, if for any $u \in \mathbb{F}^n$, $u^t A u = 0$. That is, A represents an alternating bilinear form. Note that in characteristic $\neq 2$, alternating is the same as skew-symmetric, but in characteristic 2 they differ (in characteristic 2, skew-symmetric=symmetric). The linear space of $n \times n$ alternating matrices over $\mathbb{F}$ is denoted by $\Lambda(n, \mathbb{F})$.

The $n \times n$ *identity matrix* is denoted by I_n , and when n is clear from the context, we may just write I . The *elementary matrix* $E_{i,j}$ is the matrix with the (i, j) th entry being 1, and other entries being 0. The (i, j) -th *elementary alternating matrix* is the matrix $E_{i,j} - E_{j,i}$.

Matrix tuples. We use $M(\ell \times n, \mathbb{F})^m$ to denote the linear space of m -tuples of $\ell \times n$ matrices. Boldface letters like $\mathbf{A}$ and $\mathbf{B}$ denote matrix tuples. Let $\mathbf{A} = (A_1, \dots, A_m)$, $\mathbf{B} = (B_1, \dots, B_m) \in M(\ell \times n, \mathbb{F})^m$. Given $P \in M(\ell, \mathbb{F})$ and $Q \in M(n, \mathbb{F})$, $PAQ := (PA_1Q, \dots, PA_mQ) \in M(\ell \times n, \mathbb{F})^m$. Given $R = (r_{i,j})_{i,j \in [m]} \in M(m, \mathbb{F})$, $\mathbf{A}^R := (A'_1, \dots, A'_m) \in M(\ell \times n, \mathbb{F})$ where $A'_i = \sum_{j \in [m]} r_{j,i} A_j$.

Remark 2.1. In particular, note that the coefficients in the formula of defining A'_i correspond to the entries in the i th *column* of R . While this choice is immaterial (we could have chosen the opposite convention), all of our later calculations are consistent with this convention.

Given $\mathbf{A}, \mathbf{B} \in M(n \times n, \mathbb{F})^m$, we say that $\mathbf{A}$ and $\mathbf{B}$ are *isometric*, if there exists $P \in GL(n, \mathbb{F})$, such that $P^t \mathbf{A} P = \mathbf{B}$. Finally, $\mathbf{A}$ and $\mathbf{B}$ are *pseudo-isometric* if there exist $P \in GL(n, \mathbb{F})$ and $R \in GL(m, \mathbb{F})$, such that $P^t \mathbf{A} P = \mathbf{B}^R$.

Matrix spaces. Linear subspaces of $M(\ell \times n, \mathbb{F})$ are called matrix spaces. Calligraphic letters like $\mathcal{A}$ and $\mathcal{B}$ denote matrix spaces. By a slight abuse of notation, for $\mathbf{A} \in M(\ell \times n, \mathbb{F})^m$, we use $\langle \mathbf{A} \rangle$ to denote the subspace spanned by those matrices in $\mathbf{A}$. For $\mathbf{A}, \mathbf{B} \in M(n, \mathbb{F})^m$, we say that the spaces $\langle \mathbf{A} \rangle, \langle \mathbf{B} \rangle$ are isometric iff the tuples $\mathbf{A}, \mathbf{B}$ are pseudo-isometric.

3-way arrays. Let $T(\ell \times n \times m, \mathbb{F})$ be the linear space of $\ell \times n \times m$ 3-way arrays over $\mathbb{F}$. We use the fixed-width teletypefont for 3-way arrays, like $\mathbf{A}, \mathbf{B}$, etc..

Given $\mathbf{A} \in T(\ell \times n \times m, \mathbb{F})$, we can think of $\mathbf{A}$ as a 3-dimensional table, where the (i, j, k) th entry is denoted as $A(i, j, k) \in \mathbb{F}$. We can slice $\mathbf{A}$ along one direction and obtain several matrices, which are then called slices. For example, slicing along the first coordinate, we obtain the *horizontal* slices, namely ℓ matrices $A_1, \dots, A_\ell \in M(n \times m, \mathbb{F})$, where $A_i(j, k) = A(i, j, k)$. Similarly, we also obtain the *lateral* slices by slicing along the second coordinate, and the *frontal* slices by slicing along the third coordinate.

We will often represent a 3-way array as a matrix whose entries are vectors. That is, given $\mathbf{A} \in T(\ell \times n \times m, \mathbb{F})$, we can write

$$\mathbf{A} = \begin{bmatrix} w_{1,1} & w_{1,2} & \dots & w_{1,n} \\ w_{2,1} & w_{2,2} & \dots & w_{2,n} \\ \vdots & \ddots & \ddots & \vdots \\ w_{\ell,1} & w_{\ell,2} & \dots & w_{\ell,n} \end{bmatrix},$$

where $w_{i,j} \in \mathbb{F}^m$, so that $w_{i,j}(k) = A(i, j, k)$. Note that, while $w_{i,j} \in \mathbb{F}^m$ are column vectors, in the above representation of $\mathbf{A}$, we should think of them as along the direction “orthogonal to the paper.” Following [48], we call $w_{i,j}$ the *tube fibers* of $\mathbf{A}$. Similarly, we can have the *row fibers* $v_{i,k} \in \mathbb{F}^n$ such that $v_{i,k}(j) = A(i, j, k)$, and the *column fibers* $u_{j,k} \in \mathbb{F}^\ell$ such that $u_{j,k}(i) = A(i, j, k)$.

Given $P \in M(\ell, \mathbb{F})$ and $Q \in M(n, \mathbb{F})$, let PAQ be the $\ell \times n \times m$ 3-way array whose k th frontal slice is PA_kQ . For $R = (r_{i,j}) \in GL(m, \mathbb{F})$, let A^R be the $\ell \times n \times m$ 3-way array whose k th frontal slice is $\sum_{k' \in [m]} r_{k',k} A_{k'}$. Note that these notations are consistent with the notations for matrix tuples above, when we consider the matrix tuple $A = (A_1, \dots, A_m)$ of frontal slices of A .

3 WARM UP: REDUCING MONOMIAL CODE EQUIVALENCE TO TENSOR ISOMORPHISM

The purpose of this section is to present a concrete example that illustrates what we mean by a gadget restricting to monomial subgroups. We also explain why the gadget would be viewed as a linear algebraic analogue of attaching stars in the graph setting as mentioned in Section 1.2.1.

We will give a reduction here to the TENSOR ISOMORPHISM (TI) problem, so we begin by recalling its definition:

Definition 3.1 (The d -TENSOR ISOMORPHISM problem). d -TENSOR ISOMORPHISM over a field $\mathbb{F}$ is the problem: given two d -way arrays $A = (a_{i_1, \dots, i_d})$ and $B = (b_{i_1, \dots, i_d})$, where $i_k \in [n_k]$ for $k \in [d]$, and $a_{i_1, \dots, i_d}, b_{i_1, \dots, i_d} \in \mathbb{F}$, decide whether there are $P_k \in GL(n_k, \mathbb{F})$ for $k \in [d]$, such that for all $i_1, \dots, i_d$,

$$a_{i_1, \dots, i_d} = \sum_{j_1, \dots, j_d} b_{j_1, \dots, j_d} (P_1)_{i_1, j_1} (P_2)_{i_2, j_2} \cdots (P_d)_{i_d, j_d}.$$

Let A be an $\ell \times n \times m$ 3-way array, with lateral slices $L_1, L_2, \dots, L_n$ (each an $\ell \times m$ matrix). For any vector $v \in \mathbb{F}^n$, we get an associated lateral matrix L_v , which is a linear combination of the lateral slices as given, namely $L_v := \sum_{j=1}^n v_j L_j$ (note that when $v = \vec{e}_j$ is the j -th standard basis vector, the associated lateral matrix is indeed L_j). By analogy with adjacency matrices of graphs, L_v is a natural analogue of the neighborhood of a vertex in a graph. Correspondingly, we get a notion of “degree,” which we may define as

$$\deg_A(v) := \text{rk} L_v = \text{rk} \left(\sum_{j=1}^n v_j L_j \right) = \dim \text{span} \{ L_v w : w \in \mathbb{F}^m \} = \dim \text{span} \{ u^t L_v : u \in \mathbb{F}^\ell \}.$$

The last two characterizations are analogous to the fact that the degree of a vertex v in a graph G may be defined as the number of “in-neighbors” (nonzero entries the corresponding row of the adjacency matrix) or the number of “out-neighbors” (nonzero entries in the corresponding column).

To “individualize” v , we can enlarge A with a gadget to increase $\deg_A(v)$, as in the graph case. Note that $\deg_A(v) \leq \min\{\ell, m\}$ because the lateral matrices are all of size $\ell \times m$. For notational simplicity, let us individualize $v = \vec{e}_1 = (1, 0, \dots, 0)^t$. To individualize v , we will increase its degree by $d = \min\{\ell, m\} + 1 > \max_{v \in \mathbb{F}^n} \deg_A(v)$. Extend A to a new 3-way array A_v of size $(\ell + d) \times n \times (m + d)$; in the “first” $\ell \times n \times m$ “corner”, we will have the original array A , and then we will append to it an identity matrix in one slice to increase $\deg(v)$. More specifically, the lateral slices of A_v will be

$$L'_1 = \begin{bmatrix} L_1 & 0 \\ 0 & I_d \end{bmatrix} \quad \text{and} \quad L'_j = \begin{bmatrix} L_j & 0 \\ 0 & 0 \end{bmatrix} \quad (\text{for } j > 1).$$

Now we have that $\deg_{A_v}(v) \geq d$. This almost does what we want, but now note that any vector $w = (w_1, \dots, w_n)$ with $w_1 \neq 0$ has $\deg_{A_v}(w) = \text{rk}(w_1 L'_1 + \sum_{j \geq 2} w_j L'_j) \geq d$. We can nonetheless consider this a sort of linear-algebraic individualization.

Leveraging this trick, we can then individualize an entire basis of $\mathbb{F}^n$ simultaneously, so that $d \leq \deg(v) < 2d$ for any vector v in our basis, and $\deg(v') \geq 2d$ for any nonzero v' outside the basis (not a scalar multiple of one of the basis vectors), as we do in the following result. This is also a 3-dimensional analogue of the reduction from GI to CODEEQ [54, 61, 64] (where they use Hamming weight instead of rank).

We now come to the concrete result. Given two $d \times n$ matrices A, B over $\mathbb{F}$ of rank d , the MONOMIAL CODE EQUIVALENCE problem is to decide whether there exist $Q \in GL(d, \mathbb{F})$ and a monomial matrix $P \in \text{Mon}(n, \mathbb{F}) \leq GL(n, \mathbb{F})$ (product of a diagonal matrix and a permutation matrix) such that $QAP = B$. Monomial equivalence of

linear codes is a basic notion in coding theory [12], and MONOMIAL CODE EQUIVALENCE was recently studied in the context of post-quantum cryptography [69].

Mostly for notational convenience, we make use of the following observation in the proof below:

Observation 3.2. *Two 3-tensors A, B are isomorphic iff there exists invertible matrices Q, P, R such that $QAP = B^R$.*

PROOF. With this notation, the definition of tensor isomorphism given above says that A, B are isomorphic iff there exist invertible Q', P', R such that $A = (Q'BP')^R$. Let $Q = (Q')^{-1}$, $P = (P')^{-1}$. Since the three actions (on the left, on the right, and in the third direction) commute, we have

$$\begin{aligned} A &= (Q'BP')^R \\ QA &= (BP')^R \\ QAP &= B^R. \end{aligned}$$

□

Proposition 3.3. MONOMIAL CODE EQUIVALENCE reduces to 3-TENSOR ISOMORPHISM.

PROOF. Without loss of generality we assume $d > 1$, as the problem is easily solvable when $d = 1$. We treat a $d \times n$ matrix A as a 3-way array of size $d \times n \times 1$, and then follow the outline proposed above, of individualizing the entire standard basis $\vec{e}_1, \dots, \vec{e}_n$. Since the third direction only has length 1, the maximum degree of any column is 1, so it suffices to use gadgets of rank 2. More specifically, (see Figure 1) we build a $(d + 2n) \times n \times (1 + 2n)$ 3-way array A whose lateral slices are

$$L_j = \begin{bmatrix} a_{1,j} & \mathbf{0}_{1 \times 2} & \mathbf{0}_{1 \times 2} & \cdots & \mathbf{0}_{1 \times 2} & \cdots & \mathbf{0}_{1 \times 2} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ a_{d,j} & \mathbf{0}_{1 \times 2} & \mathbf{0}_{1 \times 2} & \cdots & \mathbf{0}_{1 \times 2} & \cdots & \mathbf{0}_{1 \times 2} \\ \mathbf{0}_{2 \times 1} & \mathbf{0}_{2 \times 2} & \mathbf{0}_{2 \times 2} & \cdots & \mathbf{0}_{2 \times 2} & \cdots & \mathbf{0}_{2 \times 2} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ \mathbf{0}_{2 \times 1} & \mathbf{0}_{2 \times 2} & \mathbf{0}_{2 \times 2} & \cdots & I_2 & \cdots & \mathbf{0}_{2 \times 2} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ \mathbf{0}_{2 \times 1} & \mathbf{0}_{2 \times 2} & \mathbf{0}_{2 \times 2} & \cdots & \mathbf{0}_{2 \times 2} & \cdots & \mathbf{0}_{2 \times 2} \end{bmatrix}$$

where the I_2 block is in the j -th block of size 2 (that is, rows $d + 2(j - 1) + \{1, 2\}$ and columns $1 + 2(j - 1) + \{1, 2\}$).

It will also be useful to visualize the frontal slices of A, as follows. Here each entry of the “matrix” below is actually a $(1 + 2n)$ -dimensional vector, “coming out of the page”:

$$A = \begin{bmatrix} \tilde{a}_{1,1} & \tilde{a}_{1,2} & \cdots & \tilde{a}_{1,n} \\ \vdots & \vdots & \ddots & \vdots \\ \tilde{a}_{d,1} & \tilde{a}_{d,2} & \cdots & \tilde{a}_{d,n} \\ e_{1,1} & \mathbf{0} & \cdots & \mathbf{0} \\ e_{1,2} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & e_{2,1} & \cdots & \mathbf{0} \\ \mathbf{0} & e_{2,2} & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & e_{n,1} \\ \mathbf{0} & \mathbf{0} & \cdots & e_{n,2} \end{bmatrix}, \quad \begin{aligned} &\text{where} \\ \tilde{a}_{i,j} &= \begin{bmatrix} a_{i,j} \\ \mathbf{0}_{2n \times 1} \end{bmatrix} \in \mathbb{F}^{1+2n} \\ e_{i,j} &= \vec{e}_{1+2(i-1)+j} \in \mathbb{F}^{1+2n} \text{ for } i \in [n], j \in [2] \end{aligned}$$

and the frontal slices are

$$\begin{aligned} A_1 &= \begin{bmatrix} A \\ \mathbf{0}_{2n \times n} \end{bmatrix} \\ A_{1+2(i-1)+j} &= \vec{E}_{d+2(i-1)+j,i} \quad \text{for } i \in [n], j \in [2] \end{aligned}$$

(In A we turn the vectors $\tilde{a}_{i,j}$ and $e_{i,j}$ “on their side” so they become perpendicular to the page.)




Fig. 1. Pictorial representation of the reduction for Proposition 3.3.

We claim that A and B are monomially equivalent as codes if and only if A and B are isomorphic as 3-tensors.

($\Rightarrow$) Suppose $QADP = B$ where $Q \in \text{GL}(d, \mathbb{F})$, $D \in \text{diag}(n, \mathbb{F})$ and $P \in S_n \leq \text{GL}(n, \mathbb{F})$. Then by examining the frontal slices it is not hard to see that for $Q' = \begin{bmatrix} Q & 0 \\ 0 & (DP)^{-1} \otimes I_2 \end{bmatrix}$ (where $(DP)^{-1} \otimes I_2$ denotes a $2n \times 2n$ block matrix, where the pattern of the nonzero blocks and the scalars are governed by $(DP)^{-1}$, and each 2×2 block is either zero or a scalar multiple of I_2) we have $Q'A_1DP = B_1$ and $Q'A_{1+2(i-1)+j}DP = B_{1+2(\pi(i)-1)+j}$, where π is the permutation corresponding to P . Thus A and B are isomorphic tensors, via the isomorphism $(Q', DP, I_1 \oplus (P \otimes I_2))$, where $I_1 \oplus (P \otimes I_2)$ denotes the block-diagonal matrix $\begin{bmatrix} 1 & 0 \\ 0 & P \otimes I_2 \end{bmatrix}$.

($\Leftarrow$) Suppose there exist $Q \in \text{GL}(d + 2n, \mathbb{F})$, $P \in \text{GL}(n, \mathbb{F})$, and $R \in \text{GL}(1 + 2n, \mathbb{F})$, such that $QAP = B^R$. First, note that every lateral slice of A is of rank either 2 or 3, and the actions of Q and R do not change the ranks of the lateral slices. Furthermore, any non-trivial linear combination of more than 1 lateral slice results in a lateral matrix of rank ≥ 4 . It follows that P cannot take nontrivial linear combinations of the lateral slices, hence it must be monomial.

Now consider the frontal slices. Note that, as we assume $d > 1$, every frontal slice of QAP , except the first one, is of rank 1. Therefore, R must be of the form $\begin{bmatrix} r_{1,1} & \mathbf{0}_{1 \times (n-1)} \\ \vec{r}' & R' \end{bmatrix}$ where R' is $(n-1) \times (n-1)$. Since R is invertible, we must have $r_{1,1} \neq 0$, and the first frontal slice of B^R contains all the rows of B scaled by $r_{1,1}$ in its first d rows. The first frontal slice of QAP is a matrix that generates, by definition (and since we've shown P is monomial), a code monomially equivalent to A . Since the first frontal slices of QAP and B^R are equal, and the latter is just a scalar multiple of B_1 , we have that A and B are monomially equivalent as codes as well. $\square$

4 SEARCH-TO-DECISION REDUCTION BY RESTRICTING TO MONOMIAL GROUPS

4.1 The gadget restricting to the monomial group

In this section, we present the gadget that restricts to the monomial group in the setting of ALTERNATING MATRIX SPACE ISOMETRY. To show this, we will need the concept of monomial isometry; see Some Groups above. Recall that a matrix is monomial if, equivalently, it can be written as DP where D is a nonsingular diagonal matrix and P is a permutation matrix. We say two matrix spaces $\mathcal{A}, \mathcal{B}$ are *monomially isometric* if there is some $M \in \text{Mon}(n, \mathbb{F})$ such that $M^t \mathcal{A} M = \mathcal{B}$.

Lemma 4.1. ALTERNATING MATRIX SPACE MONOMIAL ISOMETRY *reduces to* ALTERNATING MATRIX SPACE ISOMETRY.

More specifically, there is a $\text{poly}(n, m)$ -time algorithm r taking alternating matrix tuples to alternating matrix tuples, such that for $\mathbf{A}, \mathbf{B} \in \Lambda(n, \mathbb{F})^m$, the matrix spaces $\mathcal{A} = \langle \mathbf{A} \rangle$ and $\mathcal{B} = \langle \mathbf{B} \rangle$ are monomially isometric if and only if the matrix spaces $\langle r(\mathbf{A}) \rangle$ and $\langle r(\mathbf{B}) \rangle$ are isometric.

The gadget used in Lemma 4.1 is essentially applying the gadget in Proposition 3.3 “in two directions.” Still, to prove the correctness requires some work.

PROOF. For $\mathbf{A} = (A_1, \dots, A_m) \in \Lambda(n, \mathbb{F})^m$, define $r(\mathbf{A})$ to be the alternating matrix tuple $\tilde{\mathbf{A}} = (\tilde{A}_1, \dots, \tilde{A}_{m+n^2}) \in \Lambda(n + n^2, \mathbb{F})^{m+n^2}$, where

- (1) For $k = 1, \dots, m$, $\tilde{A}_k = \begin{bmatrix} A_k & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$.
- (2) For $k = m + (i - 1)n + j$, $i \in [n]$, $j \in [n]$, $\tilde{A}_k$ is the elementary alternating matrix $E_{i, in+j} - E_{in+j, i}$.

At this point, some readers may wish to look at the large matrix in Equation 2 and/or at Figure 2.

It is clear that r can be computed in time $\tilde{O}((m + n^2)(n^2 + n)) = \text{poly}(n, m)$. Given alternating matrix tuples $\mathbf{A}, \mathbf{B}$, let $\mathcal{A}, \mathcal{B}$ be the corresponding matrix spaces they span, and let $\tilde{\mathcal{A}} = \langle r(\mathbf{A}) \rangle$ and $\tilde{\mathcal{B}} = \langle r(\mathbf{B}) \rangle$. We claim that $\mathcal{A}$ and $\mathcal{B}$ are monomially isometric if and only if $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{B}}$ are isometric.

To prove this, it will help to think of our matrix tuples $\mathbf{A}, \tilde{\mathbf{A}}$, etc. as (corresponding to) 3-way arrays, and to view these 3-way arrays from two different directions. Towards this end, write the 3-way array corresponding to $\mathbf{A}$ as

$$\mathbf{A} = \begin{bmatrix} \mathbf{0} & a_{1,2} & a_{1,3} & \dots & a_{1,n} \\ -a_{1,2} & \mathbf{0} & a_{2,3} & \dots & a_{2,n} \\ -a_{1,3} & -a_{2,3} & \mathbf{0} & \dots & a_{3,n} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ -a_{1,n} & -a_{2,n} & -a_{3,n} & \dots & \mathbf{0} \end{bmatrix},$$

where $a_{i,j}$ are vectors in $\mathbb{F}^m$ (“coming out of the page”), namely $a_{i,j}(k) = A_k(i, j)$. The frontal slices of this array are precisely the matrices $A_1, \dots, A_m$.

The 3-way array corresponding to $\tilde{\mathbf{A}} = r(\mathbf{A})$ is then the $(n+1)n \times (n+1)n \times (m+n^2)$ array:

$$\tilde{\mathbf{A}} = \begin{bmatrix} \begin{array}{cccc|cccc|cccc|cccc} 0 & \tilde{a}_{1,2} & \tilde{a}_{1,3} & \dots & \tilde{a}_{1,n} & e_{1,1} & \dots & e_{1,n} & 0 & \dots & 0 & \dots & 0 & \dots & 0 \\ -\tilde{a}_{1,2} & 0 & \tilde{a}_{2,3} & \dots & \tilde{a}_{2,n} & 0 & \dots & 0 & e_{2,1} & \dots & e_{2,n} & \dots & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ -\tilde{a}_{1,n} & -\tilde{a}_{2,n} & -\tilde{a}_{3,n} & \dots & 0 & 0 & \dots & 0 & 0 & \dots & 0 & \dots & e_{n,1} & \dots & e_{n,n} \\ -e_{1,1} & 0 & 0 & \dots & 0 & 0 & \dots & 0 & 0 & \dots & 0 & \dots & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \dots & \vdots & \vdots & \dots & \vdots & \vdots & \dots & \vdots & \dots & \vdots & \dots & \vdots \\ -e_{1,n} & 0 & 0 & \dots & 0 & 0 & \dots & 0 & 0 & \dots & 0 & \dots & 0 & \dots & 0 \\ 0 & -e_{2,1} & 0 & \dots & 0 & 0 & \dots & 0 & 0 & \dots & 0 & \dots & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \dots & \vdots & \vdots & \dots & \vdots & \vdots & \dots & \vdots & \dots & \vdots & \dots & \vdots \\ 0 & -e_{2,n} & 0 & \dots & 0 & 0 & \dots & 0 & 0 & \dots & 0 & \dots & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \dots & \vdots & \vdots & \dots & \vdots & \vdots & \dots & \vdots & \dots & \vdots & \dots & \vdots \\ 0 & 0 & 0 & \dots & -e_{n,1} & 0 & \dots & 0 & 0 & \dots & 0 & \dots & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \dots & \vdots & \vdots & \dots & \vdots & \vdots & \dots & \vdots & \dots & \vdots & \dots & \vdots \\ 0 & 0 & 0 & \dots & -e_{n,n} & 0 & \dots & 0 & 0 & \dots & 0 & \dots & 0 & \dots & 0 \end{array} \end{bmatrix}, \quad (2)$$

where $\tilde{a}_{i,j} = \begin{bmatrix} a_{i,j} \\ \mathbf{0} \end{bmatrix} \in \mathbb{F}^{m+n^2}$ (here think of the vector $a_{i,j}$ as a column vector, *not* coming out of the page; in the above array we then lay the column vector $\tilde{a}_{i,j}$ “on its side” so that it is coming out of the page), and $e_{i,j} := e_{m+(i-1)n+j} \in \mathbb{F}^{m+n^2}$, which we can equivalently write as $\begin{bmatrix} \mathbf{0}_m \\ e_i \otimes e_j \end{bmatrix}$, where we think of $e_i \otimes e_j$ here as a vector of length n^2 . Note that all the nonzero blocks besides upper-left “A” block only have nonzero entries that are strictly *further back* than the nonzero entries in the upper-left block.

The second viewpoint, which we will also use below, is to consider the lateral slices of $\tilde{\mathbf{A}}$, or equivalently, to view $\tilde{\mathbf{A}}$ from the side. When viewing $\tilde{\mathbf{A}}$ from the side, we see the $(n+1)n \times (m+n^2) \times (n+1)n$ 3-way array:

$$\tilde{\mathbf{A}}^{lat} = \begin{bmatrix} \begin{array}{cccc|cccc|cccc|cccc} \ell_{1,1} & \ell_{1,2} & \dots & \ell_{1,m} & e_{n+1} & \dots & e_{2n} & \dots & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots \\ \ell_{n,1} & \ell_{n,2} & \dots & \ell_{n,m} & 0 & \dots & 0 & \dots & e_{n^2+1} & \dots & e_{n^2+n} \\ 0 & 0 & \dots & 0 & -e_1 & \dots & 0 & \dots & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 & 0 & \dots & -e_1 & \dots & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 & 0 & \dots & 0 & \dots & -e_n & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 & 0 & \dots & 0 & \dots & 0 & \dots & -e_n \end{array} \end{bmatrix}, \quad (3)$$

where every $\ell_{i,k} \in \mathbb{F}^{n^2+n}$ has only the first n components being possibly non-zero, namely, $\ell_{i,k}(j) = A_k(i, j)$ for $i \in [n], j \in [n], k \in [m]$ and $\ell_{i,k}(j) = 0$ for any $j > n$.

(Monomial isometry of input implies isometry of output) Suppose there exist $P \in \text{Mon}(n, \mathbb{F})$ such that $\langle P^t A P \rangle = \langle B \rangle$. This happens if and only if there is an invertible matrix $Q \in \text{GL}(m, \mathbb{F})$ such that, for all i ,




Fig. 2. Pictorial representation of the reduction for Lemma 4.1.

$P^t A_i P = \sum_j Q_{ji} B_j$, or, using our shorthand notation, $P^t \mathbf{A} P = \mathbf{B}^Q$. We can construct $\tilde{P} \in \text{Mon}(n + n^2, \mathbb{F})$ and $\tilde{Q} \in \text{GL}(m + n^2, \mathbb{F})$ such that $\tilde{P}^t \tilde{\mathbf{A}} \tilde{P} = \tilde{\mathbf{B}}^{\tilde{Q}}$. In fact, we will show that we can take $\tilde{P} = \begin{bmatrix} P & \mathbf{0} \\ \mathbf{0} & P' \end{bmatrix}$ where $P' \in \text{Mon}(n^2, \mathbb{F})$, and $\tilde{Q} = \begin{bmatrix} Q & \mathbf{0} \\ \mathbf{0} & Q' \end{bmatrix}$ where $Q' \in \text{Mon}(n^2, \mathbb{F})$. It is not hard to see that this form already ensures that the first m matrices in the vector $\tilde{P}^t \tilde{\mathbf{A}} \tilde{P}$ and those of $\tilde{\mathbf{B}}^{\tilde{Q}}$ are the same, since when $\tilde{P}, \tilde{Q}$ are of this form, those first m matrices are controlled entirely by the P (resp., Q) in the upper-left block of $\tilde{P}$ (resp., $\tilde{Q}$).

The remaining question is then how to design appropriate P' and Q' to take care of the last n^2 matrices in $\tilde{\mathbf{A}}, \tilde{\mathbf{B}}$. This actually boils down to applying the following simple identity, but “in 3 dimensions:” Let P be the permutation matrix corresponding to $\sigma \in S_n$, so that $P e_i = e_{\sigma(i)}$, and $e_i^t P = e_{\sigma^{-1}(i)}^t$. Let $D = \text{diag}(\alpha_1, \dots, \alpha_n)$ be a diagonal matrix. Then

$$P^t D P = \text{diag}(\alpha_{\sigma^{-1}(1)}, \dots, \alpha_{\sigma^{-1}(n)}). \quad (4)$$

To see how Equation 4 helps in our setting, it is easier to focus attention on the lower right $n^2 \times n^2$ sub-array of $\tilde{\mathbf{A}}^{lat}$, namely:

$$M = - \begin{bmatrix} e_1 & \dots & 0 & \dots & 0 & \dots & 0 \\ \vdots & \ddots & \vdots & \dots & \vdots & \ddots & \vdots \\ 0 & \dots & e_1 & \dots & 0 & \dots & 0 \\ \hline \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & \dots & e_n & \dots & 0 \\ \hline \vdots & \ddots & \vdots & \dots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & \dots & 0 & \dots & e_n \end{bmatrix}, \quad (5)$$

The corresponding parts of the corresponding lateral slices of $(\tilde{P}^t \tilde{A} \tilde{P})^{\tilde{Q}}$ are then of the form $(P'^t M Q')^P$. Here the P in the “exponent” acts by sending the e_i entries in M to $\alpha_{\sigma(i)} e_{\sigma(i)}$ entries, where σ is the permutation supporting P and α_i is the value of the unique nonzero entry in the i -th row of P . That is, we have

$$M^P = - \begin{bmatrix} \alpha_{\sigma(1)} e_{\sigma(1)} & \dots & 0 & \dots & 0 & \dots & 0 \\ \vdots & \ddots & \vdots & \dots & \vdots & \ddots & \vdots \\ 0 & \dots & \alpha_{\sigma(1)} e_{\sigma(1)} & \dots & 0 & \dots & 0 \\ \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & \dots & \alpha_{\sigma(n)} e_{\sigma(n)} & \dots & 0 \\ \vdots & \ddots & \vdots & \dots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & \dots & 0 & \dots & \alpha_{\sigma(n)} e_{\sigma(n)} \end{bmatrix},$$

So setting $P' = \sigma \otimes I_n$, Q' the monomial matrix supported by $\sigma \otimes I_n$ with scalars being $1/\alpha_i$'s, we have $P'^t M^P Q' = M$ by Equation 4.

(Isometry of output implies monomial isometry of input) Suppose there exist $\tilde{P} \in \text{GL}(n + n^2, \mathbb{F})$ and $\tilde{Q} \in \text{GL}(m + n^2, \mathbb{F})$, such that $\tilde{P}^t \tilde{A} \tilde{P} = \tilde{B} \tilde{Q}$. The key feature of these gadgets now comes into play: consider the lateral slices of $\tilde{A}$, which are the frontal slices of A^{lat} (which may be easier to visualize by looking at Equation 3 and Figure 2). The first n lateral slices of $\tilde{A}$ and $\tilde{B}$ are of rank $\geq n$ and $< 2n$, while the other lateral slices are of rank $< n$ (in fact, they are of rank 1; note that without loss of generality we may assume $n > 1$, for the only 1×1 alternating matrix space is the zero space). Furthermore, left multiplying a lateral slice by $\tilde{P}^t$ and right multiplying it by $\tilde{Q}$ does not change its rank. However, the action of $\tilde{P}$ here is by $\tilde{P}^t \tilde{A} \tilde{P}$, and while the $\tilde{P}^t$ here corresponds to left multiplication on the lateral slices (=frontal slices of A^{lat}), the $\tilde{P}$ on the right here corresponds to taking linear combinations of the lateral slices. In other words, just as A^{lat} is the “side view” of $\tilde{A}$, $(\tilde{P}^t A^{\text{lat}} \tilde{Q})^{\tilde{P}}$ is the side view of $(\tilde{P}^t \tilde{A} \tilde{P})^{\tilde{Q}}$. Taking linear combinations of the lateral slices could, in principle, alter their rank; we will use the latter possibility to show that $\tilde{P}$ must be of a constrained form.

Write $\tilde{P} = \begin{bmatrix} P_{1,1} & P_{1,2} \\ P_{2,1} & P_{2,2} \end{bmatrix}$ where $P_{1,1}$ is of size $n \times n$. We first claim that $P_{1,2} = 0$. For if not, then in $(A^{\text{lat}})^{\tilde{P}}$ (the side view), one of the last n^2 frontal slices receives a nonzero contribution from one of the first n frontal slices of A^{lat} . Looking at the form of these slices from Equation 3, we see that any such nonzero combination will have rank $\geq n$, but this is a contradiction since the corresponding slice in B^{lat} has rank 1. Thus $P_{1,2} = 0$, and therefore $P_{1,1}$ must be invertible, since $\tilde{P}$ is.

Finally, we claim that $P_{1,1}$ has to be a monomial matrix. If not, then some frontal slice of $(A^{\text{lat}})^{\tilde{P}}$ among the first n would have a contribution from more than one of these n slices. Considering the lower-right $n^2 \times n^2$ sub-matrix of such a slice, we see that it would have rank exactly kn for some $k \geq 2$, which is again a contradiction since the first n slices of B^{lat} all have rank $< 2n$. It follows that $P_{1,1}^t A_i P_{1,1}$, $i \in [m]$, are in $\mathcal{B}$, and thus $\mathcal{A}$ and $\mathcal{B}$ are monomially isometric via $P_{1,1}$. $\square$

4.1.1 Application: reducing GRAPH ISOMORPHISM to ALTERNATING MATRIX SPACE ISOMETRY. An application of the monomial-restricting gadget is to give an immediate reduction from GRAPH ISOMORPHISM to ALTERNATING MATRIX SPACE ISOMETRY. While a reduction between these two problems is already known (cf. [35] for details), we choose to present it as an illustration of using this gadget.

Proposition 4.2. GRAPH ISOMORPHISM reduces to ALTERNATING MATRIX SPACE ISOMETRY.

PROOF. For a graph $G = ([n], E)$, let A_G be the alternating matrix tuple $A_G = (A_1, \dots, A_{|E|})$ with $A_e = E_{i,j} - E_{j,i}$ where $e = \{i, j\} \in E$, and let $\mathcal{A}_G = \langle A_G \rangle$ be the alternating matrix space spanned by that tuple. If P is a permutation

matrix giving an isomorphism between two graphs G and H , then it is easy to see that $P^t \mathcal{A}_G P = \mathcal{A}_H$, and thus the corresponding matrix spaces are isometric. The converse direction is not clear, though it is recently shown to be true in [37] with a rather intricate proof. Instead, we will provide a conceptually simpler proof, by showing that this construction gives a reduction to *monomial* isometry, and then using Lemma 4.1 to reduce to ordinary ALTERNATING MATRIX SPACE ISOMETRY.

Let us thus establish that the preceding construction gives a reduction from GI to ALTERNATING MATRIX SPACE MONOMIAL ISOMETRY. We will show that $G \cong H$ if and only if $\mathcal{A}_G$ and $\mathcal{A}_H$ are monomially isometric. The forward direction was handled above. For the converse, suppose $P^t D^t \mathcal{A}_G D P = \mathcal{A}_H$ where D is diagonal and P is a permutation matrix. We claim that in this case, P in fact gives an isomorphism from G to H . First let us establish that P alone gives an isometry between $\mathcal{A}_G$ and $\mathcal{A}_H$. Note that for any diagonal matrix $D = \text{diag}(\alpha_1, \dots, \alpha_n)$ and any elementary alternating matrix $E_{i,j} - E_{j,i}$, we have $D^t (E_{i,j} - E_{j,i}) D = \alpha_i \alpha_j (E_{i,j} - E_{j,i})$. Since $\mathcal{A}_G$ has a basis of elementary alternating matrices, the action of D on this basis is just to re-scale each basis element, and thus $D^t \mathcal{A}_G D = \mathcal{A}_G$. Thus, we have $P^t \mathcal{A}_G P = \mathcal{A}_H$.

Finally, note that $P^t (E_{i,j} - E_{j,i}) P = E_{\pi(i), \pi(j)} - E_{\pi(j), \pi(i)} = A_{\pi(e)}$, where $\pi \in S_n$ is the permutation corresponding to P , and by abuse of notation we write $\pi(e) = \pi(\{i, j\}) = \{\pi(i), \pi(j)\}$ as well. Since the elementary alternating matrices are linearly independent, and $\mathcal{A}_H$ has a basis of elementary alternating matrices, the only way for $A_{\pi(e)}$ to be in $\mathcal{A}_H$ is for it to be equal to one of the basis elements (one of the matrices in $\mathcal{A}_H$) or its negative. Since the edges are undirected, either of these two possibilities means that $\pi(e)$ must be an edge of H . In other words, $\pi(e)$ must be an edge of H . As P is invertible, we thus have that P gives an isomorphism $G \cong H$. $\square$

4.2 Search-to-decision reduction for ALTERNATING MATRIX SPACE ISOMETRY

Theorem A'. *Given an oracle deciding ALTERNATING MATRIX SPACE ISOMETRY, the task of finding an isometry between two alternating matrix spaces $\mathcal{A}, \mathcal{B} \in \Lambda(n, \mathbb{F}_q)$, if it exists, can be solved using at most $q^{O(n)}$ oracle queries each of size at most $O(n^2)$, and in time either $q^{O(n)} \cdot n! = q^{\tilde{O}(n)}$, or $q^{O(n+m)}$, where $m = \dim \mathcal{A}$.*

PROOF IDEA. The high level outline here is as follows. We proceed by induction to reduce to monomial isometry. Monomial isometry can be brute forced in time $n!(q-1)^n$, and in Prop 4.4 we show how to solve it in $q^{O(n+m)}$ time, giving the stated time bounds.

The induction is along the following lines, reminiscent of the individualization paradigm from GRAPH ISOMORPHISM. Suppose we have guessed vectors $v_1, \dots, v_i$ and a subspace V_i complementary to $\langle v_1, \dots, v_i \rangle$ such that there is an isometry $\mathcal{A} \rightarrow \mathcal{B}$ that sends $e_1 \mapsto v_1, \dots, e_i \mapsto v_i$ and $\langle e_{i+1}, \dots, e_n \rangle \mapsto V_i$. Now we want to guess $v_{i+1} \in V_i$ and a complement to v_{i+1} in V_i (that is, $V_i = \langle v_{i+1} \rangle \oplus V_{i+1}$) preserving this property. Note there are at most $q^{\dim V_i} \leq q^n$ choices for v_{i+1} and at most $q^{\dim V_i} \leq q^n$ choices for V_{i+1} (since it is a codimension-1 subspace of V_i). For each such choice of v_{i+1}, V_{i+1} , let P be an arbitrary map that sends $e_1 \mapsto v_1, \dots, e_i \mapsto v_i, e_{i+1} \mapsto v_{i+1}$, and $P(\langle e_{i+2}, \dots, e_n \rangle) = V_{i+1}$. Then v_{i+1}, V_{i+1} are valid choices iff, after replacing $\mathcal{A}$ by $P^t \mathcal{A} P$, the new $\mathcal{A}$ and $\mathcal{B}$ are isometric by an isometry that is monomial in the first i coordinates and general linear in the remaining $n-i$. To check whether this is indeed the case, we add gadgets to get 3-way arrays $\tilde{\mathcal{A}}_i, \tilde{\mathcal{B}}_i$ such that the latter two are pseudo-isometric iff $\mathcal{A}$ and $\mathcal{B}$ are isometric by an isometry that is monomial in the first i coordinates. We then feed $\tilde{\mathcal{A}}_i, \tilde{\mathcal{B}}_i$ to the decision oracle to check whether this is the case.

One of the key tricks here is guessing the complementary subspace at the same time we guess v_{i+1} . If we did not do that, at some point we would be guessing complementary subspaces of half codimension, of which there are $q^{\Theta(n^2)}$, which would have negated any asymptotic gain over a brute-force algorithm. $\square$

PROOF. We first present the gadget construction. Then based on this gadget, we present the search-to-decision reduction.

Gadget construction. Let $\mathbf{A} = (A_1, \dots, A_m)$ be an ordered linear basis of $\mathcal{A}$, and let $\mathbf{A} \in M(n \times n \times m, \mathbb{F}_q)$ be the 3-way array constructed from $\mathbf{A}$, so we can write

$$\mathbf{A} = \begin{bmatrix} \mathbf{0} & a_{1,2} & a_{1,3} & \dots & a_{1,n} \\ -a_{1,2} & \mathbf{0} & a_{2,3} & \dots & a_{2,n} \\ -a_{1,3} & -a_{2,3} & \mathbf{0} & \dots & a_{3,n} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ -a_{1,n} & -a_{2,n} & -a_{3,n} & \dots & \mathbf{0} \end{bmatrix},$$

where $a_{i,j} \in \mathbb{F}^m$, $1 \leq i < j \leq n$ thought of as a vector coming out of the page.

We first consider a 3-way array $\tilde{\mathbf{A}}_i$ constructed from $\mathbf{A}$, for any $1 \leq i \leq n-1$, as $\tilde{\mathbf{A}}_i =$

$$\begin{bmatrix} \mathbf{0} & a_{1,2} & \dots & a_{1,i} & a_{1,i+1} & \dots & a_{1,n} & -e_{1,1} & \dots & -e_{1,2n} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ -a_{1,2} & \mathbf{0} & \dots & a_{2,i} & a_{2,i+1} & \dots & a_{2,n} & \mathbf{0} & \dots & \mathbf{0} & -e_{2,1} & \dots & -e_{2,2n} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ -a_{1,i} & -a_{2,i} & \dots & \mathbf{0} & a_{i,i+1} & \dots & a_{i,n} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & -e_{i,1} & \dots & -e_{i,2n} & \mathbf{0} & \dots & \mathbf{0} \\ -a_{1,i+1} & -a_{2,i+1} & \dots & -a_{i,i+1} & \mathbf{0} & \dots & a_{i+1,n} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & -f_{1,1} & \dots & -f_{1,n} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ -a_{1,n} & -a_{2,n} & \dots & -a_{i,n} & -a_{i+1,n} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & -f_{n-i,1} & \dots & -f_{n-i,n} \\ e_{1,1} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ e_{1,2n} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & e_{2,1} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & e_{2,2n} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \dots & e_{i,1} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & e_{i,2n} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & f_{1,1} & \dots & f_{n-i,1} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & f_{1,n} & \dots & f_{n-i,n} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \end{bmatrix},$$

where $e_{j,k}$ is the $(m+2n(j-1)+k)$ th standard basis vector, and $f_{j,k}$ is the $(m+2ni+n(j-1)+k)$ th standard basis vector. A pictorial description can be seen by combining Figure 2 (for the $e_{j,k}$) and [35, Figure 3] (for the $f_{j,k}$).

We claim the following.

Claim 4.3. *If there exist invertible matrices P and Q to satisfy $(P^t \tilde{\mathbf{A}}_i P)^Q = \tilde{\mathbf{B}}_i$, then P must be in the form*

$$\begin{bmatrix} P_{1,1} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & P_{2,2} & \mathbf{0} \\ P_{3,1} & P_{3,2} & P_{3,3} \end{bmatrix}, \text{ where } P_{1,1} \text{ is a monomial matrix of size } i \times i, P_{2,2} \text{ is of size } (n-i) \times (n-i), \text{ and } P_{3,3} \text{ is of size } (2ni+n) \times (2ni+n).$$

Furthermore, there exist such P and Q if and only if $\mathbf{A}$ and $\mathbf{B}$ are isometric by a matrix of the form $\begin{bmatrix} P_{1,1} & \mathbf{0} \\ \mathbf{0} & P_{2,2} \end{bmatrix}$ where $P_{1,1}$ is a monomial matrix of size $i \times i$.

PROOF OF CLAIM. The idea here is to combine the arguments for the FGS gadget [31] as used in [35], and the monomial-restricting gadget introduced in Section 4.1. In fact, we will see that these two gadgets can be combined

seamlessly into the above construction, and the claim follows immediately from the aforementioned arguments. Nonetheless, for completeness, we spell out the details.

Write

$$P = \begin{bmatrix} P_{1,1} & P_{1,2} & P_{1,3} \\ P_{2,1} & P_{2,2} & P_{2,3} \\ P_{3,1} & P_{3,2} & P_{3,3} \end{bmatrix}$$

where $P_{1,1}$ is $i \times i$, $P_{2,2}$ is $(n - i) \times (n - i)$ and $P_{3,3}$ is $(2ni + n) \times (2ni + n)$.

First, we focus on the lateral slices. Note that the lateral slice of $(P^t A_i P)^Q$ are the frontal slices of $(P^t A_i^{lat} Q)^P$. Thus, the P in the “exponent” here is taking a (monomial) linear combination of the lateral slices. As the ranks of the frontal slices of $(P^t A_i^{lat} Q)$ are the same as the ranks of the frontal slices of $\tilde{A}_i^{lat}$ (=the lateral slices of $\tilde{A}_i$), we now consider their ranks. We have:

- The first i lateral slices have rank in $[2n, 3n)$. They are at least rank $2n$ because of the identity gadgets in the lower blocks. There is at most an additional rank $n - 1$ because of the entries in the first n rows. Note that this is $n - 1$ rather than n because the tube fibers (coming out of the page) along the diagonal are $\mathbf{0}$ in the upper-left $n \times n$ sub-array, giving an entire row of zeros in the lateral slice.
- The next $n - i$ lateral slices have rank in $[n, 2n)$. The lower bound of n comes from the identity gadget in the bottom-most block, and the additional $\leq n - 1$ comes from the first n rows, as in the previous case.
- Of the remaining lateral slices, the first $2ni$ of these have rank 1 (coming from the $-e_{i,j}$ in the upper-most block), and the remaining n lateral slices have rank exactly $n - i \leq n - 1$ (since $i \geq 1$) coming from the identity gadgets in the rightmost block of $\tilde{A}_i$. However, all we will need is that these remaining $2ni + n$ slices have rank in $[1, n)$.

Next we consider what happens when we take linear combinations of the lateral slices. Recall from above that P governs the linear combination of the lateral slices of $(P^t A_i^{lat} Q)^P$. When we say a linear combination “involves” a slice, we mean that slice occurs in the linear combination with nonzero coefficient.

- If a linear combination involves 1 or more of the first i lateral slices, it has rank at least $2n$ because of the identity block coming from the $e_{i,j}$. Since the only lateral slices of B_i that have rank $\geq 2n$ are the first i , this tells us that $P_{1,2} = P_{1,3} = \mathbf{0}$. Since P is invertible, this further implies that $P_{1,1}$ must be invertible.
- If a linear combination involves 2 or more of the first i lateral slices, it has rank at least $4n$, because of the identity blocks coming from the $e_{i,j}$ in the description of A_i above. Since there are no lateral slices of rank $\geq 3n$ in B_i , this tells us that $P_{1,1}$ has at most one nonzero entry per column. Since $P_{1,1}$ is invertible by the above, we have that $P_{1,1}$ is a monomial matrix.
- If a linear combination involves at least one of the first i lateral slices and at least one of the next $n - i$ lateral slices, it has rank at least $3n$: $2n$ coming from the identity gadget among the $e_{i,j}$, and another n coming from the identity gadget among the $f_{i,j}$. These two add because they are identity matrices on disjoint sets of columns in the lateral slice. Since all lateral slices of B_i have rank strictly less than $3n$, this tells us that $P_{2,1} = \mathbf{0}$.
- Finally, because the last $2ni + n$ lateral slices have rank strictly less than n , but any linear combination involving at least one of the lateral slices $i + 1, i + 2, \dots, n$ has rank $\geq n$, we have that $P_{2,3} = \mathbf{0}$ as well.

This completes the proof of the first part of the claim.

For the “furthermore,” the $(\Rightarrow)$ direction is straightforward: after observing that P has to be of the above form, we can easily verify that $\begin{bmatrix} P_{1,1} & \mathbf{0} \\ \mathbf{0} & P_{2,2} \end{bmatrix}$ is an isometry from A to B , where $P_{1,1}$ is monomial.

For ($\Leftarrow$) direction of the “furthermore,” starting from $\begin{bmatrix} P_{1,1} & \mathbf{0} \\ \mathbf{0} & P_{2,2} \end{bmatrix}$ and $Q_{1,1} \in \text{GL}(m, \mathbb{F})$, we need to design $P_{3,3} \in \text{GL}(2ni + n, \mathbb{F})$ and $Q_{2,2} \in \text{GL}(2ni + n(n-i), \mathbb{F})$ such that letting $P = \begin{bmatrix} P_{1,1} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & P_{2,2} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & P_{3,3} \end{bmatrix}$ and $Q = \begin{bmatrix} Q_{1,1} & \mathbf{0} \\ \mathbf{0} & Q_{2,2} \end{bmatrix}$, we have $P^t \tilde{A}_i P = \tilde{B}_i^Q$. For this part of the argument, we can treat the $e_{i,j}$ gadgets and the $f_{i,j}$ gadgets independently. That is, we take $P_{3,3} = \begin{bmatrix} P_{3,3,1} & \mathbf{0} \\ \mathbf{0} & P_{3,3,2} \end{bmatrix}$ and $Q_{2,2} = \begin{bmatrix} Q_{2,2,1} & \mathbf{0} \\ \mathbf{0} & Q_{2,2,2} \end{bmatrix}$, where $P_{3,3,1}$ and $Q_{2,2,1}$ are $2ni \times 2ni$, $P_{3,3,2}$ is $n \times n$ and $Q_{2,2,2}$ is $n(n-i) \times n(n-i)$. Then $P_{3,3,1}$ and $Q_{2,2,1}$ are the same as in the “Monomial isometry implies isometry” part of the proof of Lemma 4.1 (where the same “ $e_{i,j}$ ” gadgets are used), and $P_{3,3,2}$ and $Q_{2,2,2}$ are the matrices that come from the “only if” direction of [35, Proposition 3.3] (where the same “ $f_{i,j}$ ” gadgets are used). $\square$

The search-to-decision reduction. Given these preparations, we now present the search-to-decision reduction for ALTERNATING MATRIX SPACE ISOMETRY. Recall that this requires us to use the decision oracle $\mathcal{O}$ to compute an explicit isometry transformation $P \in \text{GL}(n, q)$, if $\mathcal{A}$ and $\mathcal{B}$ are indeed isometric. Think of P as sending the standard basis $(\vec{e}_1, \dots, \vec{e}_n)$ to another basis $(v_1, \dots, v_n)$, where $\vec{e}_i$ and v_i are in $\mathbb{F}_q^n$.

In the first step, we guess v_1 , the image of $\vec{e}_1$, and a complement subspace of $\langle v_1 \rangle$, at the cost of $q^{O(n)}$. For each such guess, let P_1 be the matrix which sends $\vec{e}_1 \mapsto v_1$ and sends $\langle \vec{e}_2, \dots, \vec{e}_n \rangle$ to the chosen complementary subspace arbitrarily. We apply P_1 to A , and still call the resulting 3-way array A in the following. Then construct $\tilde{A}_1$ and $\tilde{B}_1$, and feed these two instances to the oracle $\mathcal{O}$. Note that, since $P_{1,1}$ (using notation as above) must be monomial, any equivalence between $\tilde{A}_1$ and $\tilde{B}_1$ must preserve our choice of v_1 up to scale. Thus, clearly, if A and B are indeed isometric and we guess the correct image of e_1 , then the oracle $\mathcal{O}$ will return yes (and conversely).

In the second step, we guess v_2 , the image of $\vec{e}_2$, and a complement subspace of $\langle v_2 \rangle$ within $\langle \vec{e}_2, \dots, \vec{e}_n \rangle$, at the cost of $q^{O(n)}$. Note here that the previous step guarantees that there is an isometry respecting the direct sum decomposition $\langle v_1 \rangle \oplus \langle \vec{e}_2, \dots, \vec{e}_n \rangle$, so we need only search for a complement of v_2 within $\langle \vec{e}_2, \dots, \vec{e}_n \rangle$, and *not* a more general complement of $\langle v_1, v_2 \rangle$ in all of $\mathbb{F}_q^n$. This is crucial for the runtime, as at the $n/2$ step, the latter strategy would result in searching through $q^{\Theta(n^2)}$ possibilities.

For each such guess, we apply the corresponding transformation to A (and again call the resulting 3-way array A). Then construct $\tilde{A}_2$ and $\tilde{B}_2$, and feed these two instances to the oracle $\mathcal{O}$. Clearly, if $\mathcal{A}$ and $\mathcal{B}$ are indeed isometric and we guess the correct image of $\vec{e}_2$ (and $\vec{e}_1$ from the previous step), then the oracle $\mathcal{O}$ will return yes. However, there is a small caveat here, namely we may guess some image of e_2 , such that $\mathcal{A}$ and $\mathcal{B}$ are actually isometric by some matrix P of the form $\begin{bmatrix} P_{1,1} & \mathbf{0} \\ \mathbf{0} & P_{2,2} \end{bmatrix}$ where $P_{1,1}$ is a monomial matrix of size 2 (instead of the more desired diagonal matrix). But this is fine, as it still ensures $P_{1,1}$ to be monomial, which is the key property to keep. This means that our choices of $\{v_1, v_2\}$ is correct as a set up to scaling, so we proceed.

In general, in the i th step, we maintain the property that $\mathcal{A}$ and $\mathcal{B}$ are isometric by some $P = \begin{bmatrix} P_{1,1} & \mathbf{0} \\ \mathbf{0} & P_{2,2} \end{bmatrix}$ where $P_{1,1}$ is a monomial matrix of size $(i-1) \times (i-1)$. We guess v_i , the image of $\vec{e}_i$ in $\langle \vec{e}_i, \dots, \vec{e}_n \rangle$, and a complement subspace of $\langle v_i \rangle$ within $\langle \vec{e}_i, \dots, \vec{e}_n \rangle$. This cost is $q^{O(n)}$. For each such guess, we apply the corresponding transformation to A (and call the resulting 3-way array A). Then construct $\tilde{A}_i$ and $\tilde{B}_i$, and feed these two instances to the oracle $\mathcal{O}$. Once we guess correctly, we ensure that $\mathcal{A}$ and $\mathcal{B}$ are isometric by $P = \begin{bmatrix} P_{1,1} & \mathbf{0} \\ \mathbf{0} & P_{2,2} \end{bmatrix}$ where $P_{1,1}$ is a monomial matrix of size $i \times i$.

So after the $(n-1)$ th step, we know that $\mathcal{A}$ and $\mathcal{B}$ are isometric by a monomial transformation. As the number of all monomial transformations is $(q-1)^n \cdot n! \leq q^n \cdot 2^{n \log n} = q^{\tilde{O}(n)}$, we can enumerate all monomial

transformations and check correspondingly. This gives an algorithm in time $q^{\tilde{O}(n)}$. By resorting to Prop. 4.4 which solves ALTERNATING MATRIX SPACE MONOMIAL ISOMETRY in time $q^{O(n+m)}$, we have an algorithm in time $q^{O(n+m)}$.

Note that all the instances we feed into the oracle $\mathcal{O}$ are of size $O(n^2)$. This concludes the proof. $\square$

4.3 A simply-exponential algorithm for monomial isometry of alternating matrix spaces

We now state the algorithm for monomial isometry used in Theorem A'.

Proposition 4.4. *Let $\mathcal{A}, \mathcal{B} \leq \Lambda(n, q)$ be m -dimensional. Then there exists a $q^{O(n+m)}$ -time algorithm that decides whether $\mathcal{A}$ and $\mathcal{B}$ are monomially isometric, and if so, computes an explicit monomial isometry.*

PROOF. Let $\mathcal{A}, \mathcal{B} \leq \Lambda(n, q)$ be two m -dimensional alternating matrix spaces. Clearly, by incurring a multiplicative factor of q^n , we can reduce to the problem of testing whether $\mathcal{A}$ and $\mathcal{B}$ are permutationally isometric, i.e. whether there exists a permutation matrix $T \in \text{GL}(n, q)$, such that $T^t \mathcal{A} T = \mathcal{B}$. We will solve this problem in time $2^{O(n)} \cdot q^{O(m)}$. This would give an algorithm with total running time $q^n \cdot 2^{O(n)} \cdot q^{O(m)} = q^{O(n+m)}$. The basic idea of the algorithm comes from Luks's dynamic programming technique for HYPERGRAPH ISOMORPHISM [55].

Reducing to a generalized linear code equivalence problem. Suppose $\mathcal{A} = \langle A_1, \dots, A_m \rangle$, and $\mathcal{B} = \langle B_1, \dots, B_m \rangle$. Let A and B be the $n \times n \times m$ 3-way arrays formed by the given bases of $\mathcal{A}$ and $\mathcal{B}$. The group $S_n \times \text{GL}(m, q)$ acts naturally on the set of such 3-way arrays as follows: $(\pi, Q) \cdot A = (P_\pi A P_\pi^T)^Q$, where P_π is the permutation matrix corresponding to π . The action of $\text{GL}(m, q)$ here corresponds to changing basis within a subspace, and thus one sees that two such 3-way arrays are in the same orbit of this action if and only if the corresponding matrix spaces are permutationally isometric. For this proof, we introduce the notation $\text{Permlsom}(A, B)$ for the coset in $S_n \times \text{GL}(m, q)$ that sends A to B under this action.

For $S \subseteq [n]$ let A_S denote the $n \times n \times m$ 3-way array that agrees with A on indices (i, j, k) whenever both i and j are in S , and is zero outside of this region (in particular, if $|S| = s$, then the nonzero region in A_S has size $s \times s \times m$). Similarly for B_S . For two sets $S, T \subseteq [n]$, let $\text{Permlsom}_{S \rightarrow T}(A, B)$ denote the coset in $S_n \times \text{GL}(m, q)$ of permutational isometries (π, Q) that send A_S to B_T and such that $\pi(S) = T$.

Our goal is to compute $\text{Permlsom}(A, B)$. Note that $\text{Permlsom}(A, B) = \text{Permlsom}_{[n] \rightarrow [n]}(A, B)$. We will show how to compute $\text{Permlsom}(A, B)$ by inductively computing $\text{Permlsom}_{S \rightarrow T}(A, B)$ for all subsets $S, T \subseteq [n]$. (If we wanted, we could save a factor of 2^n in the runtime by only computing this $\text{Permlsom}_{[s] \rightarrow [s]}$ for all $s = 0, \dots, n$ and all subsets T , but as this is not the dominant term in the runtime, we compute $\text{Permlsom}_{S \rightarrow T}$ for all subsets S, T , which makes the presentation more symmetric in terms of A and B .)

Our base case is $S = T = \emptyset$. In this case we have that both A_S and B_T are the all-zeros arrays, and since all permutations map the empty set to itself, we have $\text{Permlsom}_{S \rightarrow T}(A, B) = S_n \times \text{GL}(m, q)$.

Now inductively suppose we have computed $\text{Permlsom}_{S \rightarrow T}(A, B)$ for all sets S and T of size $|S| = |T| = k-1 \geq 0$. We show how to compute the same for all sets of size k . Let $S, T \subseteq [n]$ be two sets of size k . Let $S = \{s_1, \dots, s_k\}$ and $S' = \{s_1, \dots, s_{k-1}\} = S \setminus \{s_k\}$. Any $(\pi, Q) \in \text{Permlsom}_{S \rightarrow T}(A, B)$ must send S' to some $T' \subset T$ of size $k-1$, so we must have $(\pi, Q) \in \text{Permlsom}_{S' \rightarrow T'}(A, B)$, which has already been computed. Let $t_k = \pi(s_k)$. On the other hand, for $(\pi, Q) \in \text{Permlsom}_{S' \rightarrow T'}(A, B)$ to be in $\text{Permlsom}_{S \rightarrow T}(A, B)$, (π, Q) needs to send the s_k -th horizontal slice of A_S to the t_k -th horizontal slice of B_T . (The same is required of the lateral slices, but this will follow automatically because frontal slices are alternating matrices.)

Let $\text{CodeEq}_{s_k, t_k}(A, B)$ denote the set of (π, Q) that send the s_k -th horizontal slice of A to the t_k -th horizontal slice of B , that is, $\pi(s_k) = t_k$ and $B(t_k, \pi(i), \ell) = \sum_{\ell'} Q_{\ell', \ell} A(s_k, i, \ell')$ for all $i \in [n], \ell \in [m]$. Then the previous paragraph can be summarized in the following equation

$$\text{Permlsom}_{S \rightarrow T}(A, B) = \bigcup_{t_k \in T} \left(\text{Permlsom}_{S' \rightarrow (T \setminus \{t_k\})}(A, B) \cap \text{CodeEq}_{s_k, t_k}(A_S, B_T) \right).$$

If we treat $\text{GL}(m, q)$ as a permutation group on q^m elements, then the entire group $S_n \times \text{GL}(m, q)$ is a permutation group on domain size nq^m . With this view, if we could compute $\text{PermIsom}_{S' \rightarrow (T \setminus \{t_k\})}(A, B) \cap \text{CodeEq}_{s_k, t_k}(A_S, B_T)$ in time $q^{O(n+m)}$, then the above equation can be computed in its entirety in time $q^{O(n+m)}$. Since the number of entries in the dynamic programming table is 2^{2n} , the total runtime will be $q^{O(n+m)}$, as claimed. The remainder of the proof shows how to compute $\text{PermIsom}_{S' \rightarrow (T \setminus \{t_k\})}(A, B) \cap \text{CodeEq}_{s_k, t_k}(A_S, B_T)$ in time $q^{O(n+m)}$.

Solving the generalized linear code equivalence problem. In fact, we will show that the following slightly more general problem can be solved in the desired time bound.

Problem 4.4, a generalization of LINEAR CODE EQUIVALENCE

Input: Elements $\rho_0, \rho_1, \dots, \rho_k \in S_n \times \text{GL}(m, q)$, two $n \times m$ matrices A, B over $\mathbb{F}_q$, and two indices $s, t \in [n]$

Output: Let $G = \langle \rho_1, \dots, \rho_k \rangle$. The output is the subcoset of $S_n \times \text{GL}(m, q)$ consisting of pairs $(\pi, Q) \in \rho_0 G$ such that $\pi(s) = t$ and $P_\pi A Q = B$.

Here the subcoset in the output is specified by a single element together with a generating set of the corresponding subgroup (the same way the subcoset is represented in the input). In the application above, we apply this problem with A being the s_k -th slice of A_S , B being the t_k -th slice of B_T , $s = s_k$, $t = t_k$, and the subcoset $\rho_0 G = \text{PermIsom}_{S' \rightarrow (T \setminus \{t_k\})}(A, B)$.

We solve Problem 4.4 again by a dynamic programming algorithm as follows. For $R, R' \subseteq [n]$ of size r , A_R denotes the $n \times m$ matrix that agrees with A in rows indexed by R , and is zero in all other rows; similarly for $B_{R'}$. Let $\text{CodeEq}_{R \rightarrow R'}^{s \rightarrow t, \rho_0 G}(A, B)$ denote the subcoset of $\rho_0 G$ consisting of those (π, Q) such that $\pi(s) = t$, $\pi(R) = R'$, and $P_\pi A_R Q = B_{R'}$. Here the information in the superscript is part of the input and will not change throughout the recursion, whereas the information the subscript will be inducted on.

The base case is $R = R' = \emptyset$, for which we have $\text{CodeEq}_{\emptyset \rightarrow \emptyset}^{s \rightarrow t, \rho_0 G}(A, B) = \{(\pi, Q) \in \rho_0 G : \pi(s) = t\}$. As above, if we view $S_n \times \text{GL}(m, q)$ as a permutation group on a set of size nq^m , then this is simply computing an element transporter inside a subcoset of a permutation group, which can be done in time $(nq^m)^{O(1)}$ [54].

Suppose inductively we have computed $\text{CodeEq}_{R \rightarrow R'}^{s \rightarrow t, \rho_0 G}(A, B)$ for all sets R, R' of size $r - 1 \geq 0$. We will show how to compute the same for all sets R, R' of size r . Fix $r_0 \in R$. For $r_0, r'_0 \in [n]$ let X_{r_0, r'_0} be the subcoset of S_n that sends r_0 to r'_0 , and for $u, v \in \mathbb{F}_q^m$ let $Y_{u, v}$ be the subcoset of $\text{GL}(m, q)$ that sends u to v . By slight abuse of notation, let A_{r_0} denote the r_0 -th row of A and $B_{r'_0}$ denote the r'_0 -th row of B .

Then, similar to the reasoning above, we have that any (π, Q) we seek must send r_0 to an element of R' , say r'_0 , and we seek the pairs $(\pi, Q) \in \text{CodeEq}_{(R \setminus \{r_0\}) \rightarrow (R' \setminus \{r'_0\})}^{s \rightarrow t, \rho_0 G}(A, B)$ such that $\pi(r_0) = r'_0$ and $A_{r_0} Q^T = B_{r'_0}$. Taking the union over all choices of $r'_0 \in R'$, we thus get the equation:

$$\text{CodeEq}_{R \rightarrow R'}^{s \rightarrow t, \rho_0 G}(A, B) = \bigcup_{r'_0 \in R'} \left(\text{CodeEq}_{(R \setminus \{r_0\}) \rightarrow (R' \setminus \{r'_0\})}^{s \rightarrow t, \rho_0 G}(A, B) \cap (X_{r_0, r'_0} \times Y_{A_{r_0}, B_{r'_0}}) \right). \quad (6)$$

Finally, we show how to efficiently compute the intersection in parentheses in the preceding equation. Let $\sigma H = \text{CodeEq}_{(R \setminus \{r_0\}) \rightarrow (R' \setminus \{r'_0\})}^{s \rightarrow t, \rho_0 G}(A, B)$. We have that $(\pi, Q) \in (\sigma H) \cap (X_{r_0, r'_0} \times Y_{A_{r_0}, B_{r'_0}})$ iff

$$\pi(r_0) = r'_0 \text{ and } A_{r_0} Q^T = B_{r'_0}.$$

Write $\sigma = (\pi_0, Q_0)$. Then we have $\pi = \pi_0 \pi'$ and $Q = Q_0 Q'$ for some $\pi' \in S_n, Q' \in \text{GL}(m, q)$, and the preceding condition is the same as

$$\pi'(r_0) = \pi_0^{-1}(r'_0) \text{ and } A_{r_0}(Q')^T = B_{r'_0}(Q_0')^T. \quad (7)$$

Since r_0, r'_0, π_0, Q_0 are all fixed, the subcoset of H consisting of (π', Q') satisfying (7) is a pointwise transporter in the permutation group $H \leq S_n \times \text{GL}(m, q)$ acting on a domain of size nq^m , which can thus be computed in time

$(nq^m)^{O(1)}$. Thus the intersection in parentheses in (6) can be computed in the same time bound. The union of subcosets can similarly be computed in time $(nq^m)^{O(1)}$ with standard permutation group machinery, and thus all of (6) can be. Again, the dynamic programming table here has size 2^{2n} , so the total runtime of this procedure is $2^{2n}(nq^m)^{O(1)} = 2^{O(n)}q^{O(m)} \leq q^{O(n+m)}$, as claimed. This completes the proof. $\square$

5 COUNTING-TO-DECISION REDUCTION BY RESTRICTING TO DIAGONAL GROUPS

In this section, we devise a gadget to achieve the restriction to the group of diagonal matrices, and use it to do the counting to decision reduction for ALTERNATING MATRIX SPACE ISOMETRY.

5.1 Describing the gadget

Let $\mathcal{A} \leq \Lambda(n, q)$ be an alternating matrix space, and let $\mathbf{A} = (A_1, \dots, A_m) \in \Lambda(n, q)^m$ be an ordered linear basis of $\mathcal{A}$. Let $\mathbf{A} \in \mathbf{T}(n \times n \times m, \mathbb{F}_q)$ be the 3-way array constructed from $\mathbf{A}$, i.e. the i th frontal slice of $\mathbf{A}$ is A_i .

We shall assume n is larger than some constant, and $q = n^{\Omega(1)}$ throughout the remainder of this section.

The form of the gadget. To describe the gadget, it is easier to view $\mathbf{A}$ from the lateral viewpoint. That is, for $i \in [n]$, let $C_i = [A_1 e_i, \dots, A_m e_i] \in M(n \times m, q)$. Let $\mathbf{C} = (C_1, \dots, C_n) \in M(n \times m, q)^n$. Then construct $\mathbf{C}' = (C'_1, \dots, C'_n)$, $C'_i = \begin{bmatrix} C_i & 0 \\ 0 & G_i \end{bmatrix}$, where G_i is of size $6n \times 4n^2$. For $i \in [n]$, $G_i = \begin{bmatrix} 0 & \dots & 0 & H_i & 0 & \dots & 0 \end{bmatrix}$, where H_i is of size $6n \times 4n$ in the i th block, and 0 denotes an all-zero matrix of size $6n \times 4n$. The H_i will be described below.

After the above step, we obtain a 3-way array $\mathbf{C} \in \mathbf{T}(7n \times n \times (m + 4n^2), \mathbb{F})$. The frontal slices of $\mathbf{C}$ are matrices of size $7n \times n$. To preserve the alternating structure, we need to do the following. Let the first n horizontal slices of $\mathbf{C}$ them be $\mathbf{B} = (B_1, \dots, B_n) \in M(n \times (m + 4n^2), \mathbb{F})$. Note that $B_i = [C_i, 0]$, where $C_i \in M(n \times m, \mathbb{F})$ was defined in the paragraph above. Then set $\mathbf{B}' = (B'_1, \dots, B'_n)$, $B'_i = \begin{bmatrix} C_i & 0 \\ 0 & -G_i \end{bmatrix}$, where $-G_i$ is of size $6n \times 4n^2$ as defined in the

above paragraph. Let $\mathbf{D}$ be one of the rest $6n$ horizontal slices of $\mathbf{C}$. Then we set $\mathbf{D}' = \begin{bmatrix} \mathbf{D} \\ 0 \end{bmatrix}$ where 0 denotes a size $6n \times (m + 4n^2)$ all-zero matrix. After the above operations, we obtain a 3-way array $\tilde{\mathbf{A}}$ of size $7n \times 7n \times (m + 4n^2)$, whose frontal slices are alternating matrices.

To summarise, from the frontal viewpoint of looking at $\mathbf{A}$, G_i 's are inserted, vertically, below and behind $\mathbf{A}$. So to preserve the alternating structure, $-G_i$'s also need to be inserted, horizontally, on the right and behind $\mathbf{A}$. We therefore get $\tilde{\mathbf{A}}$, which is of size $7n \times 7n \times (m + 4n^2)$.

Fact 5.1. *Every lateral slice of $\tilde{\mathbf{A}}$ is of rank $\leq 5n$.*

PROOF. The first n lateral slices of $\tilde{\mathbf{A}}$ are of the following form: $C'_i = \begin{bmatrix} C_i & 0 \\ 0 & G_i \end{bmatrix}$, where G_i is of size $6n \times 4n^2$. For $i \in [n]$, $G_i = \begin{bmatrix} 0 & \dots & 0 & H_i & 0 & \dots & 0 \end{bmatrix}$, where H_i is of size $6n \times 4n$ in the i th block. So $\text{rank}(C'_i) = \text{rank}(C_i) + \text{rank}(G_i) \leq n + 4n = 5n$.

The last $6n$ lateral slices of $\tilde{\mathbf{A}}$ are of the form $D_i = \begin{bmatrix} 0 & K_i \\ 0 & 0 \end{bmatrix}$ where K_i is of size $n \times 4n^2$. So $\text{rank}(D_i) = \text{rank}(K_i) \leq n$. $\square$

Remark 5.2. In the above, we attached G_i , $i \in [n]$, to each vertical slice. (And therefore, we attached $-G_i$ to each horizontal slice.) Sometimes, we may only attach G_i to the first k vertical slices. (And therefore, we only attach $-G_i$ to the first k horizontal slice.) In this case, the resulting $\tilde{\mathbf{A}}$ is of size $7n \times 7n \times (m + 4nk)$.

Conditions imposed on the H_i 's. Of course, the key to the construction above lies in the properties of the H_i 's.

Definition 5.3. Let $H_1, \dots, H_n \in M(6n \times 4n, q)$, and let $V_i \leq \mathbb{F}_q^{6n}$ be the subspace spanned by the columns of H_i . We say that the tuple $(H_1, \dots, H_n)$ is *rigid*, if the following conditions are satisfied.

- (1) For any $i \in [n]$, $\text{rk}(H_i) = \dim(V_i) = 4n$.
- (2) For any $i, j \in [n]$, $i \neq j$, $\text{rk}([H_i H_j]) = \dim(V_i \cup V_j) = 6n$.
- (3) For any $(i_1, i_2, i_3, i_4, i_5, i_6) \in [n]^6$ and $(j_1, j_2, j_3, j_4, j_5, j_6) \in [n]^6$, such that $|\{i_1, \dots, i_6\} \cup \{j_1, \dots, j_6\}| = 12$, i.e. i_k and j_ℓ all different, the coset $C = \{T \in \text{GL}(6n, q) : \forall k \in [6], T(V_{i_k}) = V_{j_k}\}$ is empty. Note that for any $i \in [n]$, $T(V_i)$ is spanned by the columns of TH_i .
- (4) For any $(i_1, i_2, i_3, i_4, i_5, i_6) \in [n]^6$, i_k all different, the group $S = \{T \in \text{GL}(6n, q) : \forall k \in [6], T(V_{i_k}) = V_{i_k}\}$ consists of only of scalar matrices.

Remark 5.4. Given $H_1, \dots, H_n \in M(6n \times 4n, q)$, whether $(H_1, \dots, H_n)$ is rigid can be verified in polynomial time as follows.

Conditions (1) and (2) are easily verified in deterministic polynomial time.

For condition (3), it can be formulated as a linear algebraic problem as follows. Let X be a $6n \times 6n$ variable matrix, so its entries are formal variables. Similarly define Y_k , $k \in [6]$, to be $4n \times 4n$ variable matrices. Then the entries of the matrix XH_{i_k} are linear forms in the variables in X . Similarly, the entries of the matrix $H_{j_k}Y_k$ are linear forms in the variables in Y_k . Equating $XH_{i_k} = H_{j_k}Y_k$, we get $4n \cdot 6n$ linear equations. Solving these linear equations, we get a linear subspace of $\mathbb{F}_q^{(6n)^2 + 6 \cdot (4n)^2}$. The question is then whether this subspace contains $(T, R_1, \dots, R_6)$ where $T \in \text{GL}(6n, q)$ and $R_i \in \text{GL}(4n, q)$. This is an instance of the symbolic determinant identity testing (SDIT) problem, so it admits a randomized efficient algorithm when $q = n^{\Omega(1)}$.

In fact, this instance of SDIT problem can be solved in deterministic polynomial time. For this let us also check out condition (4). Here, let X and Y_i be from above, and set up the equations $XH_{i_k} = H_{i_k}Y_k$. Solve the linear equations to get a subspace of $\mathbb{F}_q^{(6n)^2 + 6 \cdot (4n)^2}$. This subspace turns out to be an algebra under the natural multiplications. Indeed, if $AH_{i_k} = H_{i_k}B_k$ and $A'H_{i_k} = H_{i_k}B'_k$, then $AA'H_{i_k} = H_{i_k}B_kB'_k$. Computing the unit group in a matrix algebra can be solved by a polynomial-time Las Vegas algorithm by [17]. Given the unit group, whether it consists of only scalar matrices can be verified easily in deterministic polynomial time.

Then the linear space in condition (3) is a module over the algebra defined in the last paragraph. Because of this structure, the SDIT problem for such instances can be solved in deterministic polynomial time [15, 20, 40].

5.2 Construction and properties of the gadget

The following three propositions reveal the construction and functions of the gadget described above.

First about the construction. Instead of constructing the above H_i 's explicitly in a deterministic way, we shall show that random choices suffice.

Proposition 5.5. Suppose the entries of $H_i \in M(6n \times 4n, q)$, $i \in [n]$, are sampled uniformly and independently at random from $\mathbb{F}_q$. Then $(H_1, \dots, H_n)$ is rigid as defined in Definition 5.3 with probability $\geq 1 - \frac{n^{O(1)}}{q^{\Omega(1)}}$.

Second about the functionality. The following proposition formally explains this.

Proposition 5.6. Suppose A and B are two 3-tensors constructed from ordered bases of m -dimensional alternating matrix spaces $\mathcal{A}, \mathcal{B} \leq \Lambda(n, q)$. Let $\tilde{A}$ and $\tilde{B}$ be constructed as above, and let $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{B}}$ be the alternating matrix spaces spanned by the frontal slices of $\tilde{A}$ and $\tilde{B}$, respectively. Then $\mathcal{A}$ and $\mathcal{B}$ are isometric via a diagonal matrix if and only if $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{B}}$ are isometric.

Finally we shall use this gadget to achieve a counting-to-decision reduction for ALTERNATING MATRIX SPACE ISOMETRY. Formally, we have the following.

Proposition 5.7. *Suppose we are given $\mathcal{A}, \mathcal{B} \leq \Lambda(n, q)$ and a decision oracle for ALTERNATING MATRIX SPACE ISOMETRY. Then there exists a Las Vegas randomized algorithm that computes the number of isometries from $\mathcal{A}$ to $\mathcal{B}$ in time $q^{O(n)}$.*

The next three subsections are devoted to the proofs of Propositions 5.5 (Section 5.2.3), 5.6 (Section 5.2.1), and 5.7 (Section 5.2.2). Note that, because the proof of Proposition 5.5 is more complicated compared to the other two, we postpone it to the last.

Remark 5.8. In fact, we expect that this construction works even for small finite fields. The bottleneck lies in Proposition 5.5. If the probability $\frac{n^{O(1)}}{q^{O(1)}}$ could be improved to $\frac{n^{O(1)}}{q^{O(n)}}$, then we would be done. We believe it possible to utilize the structure of invariant subspaces under matrix actions over $\mathbb{F}_q$ to achieve this. However, we expect that the calculations will be tedious and heavy, so we hope to leave this to a future work.

5.2.1 Restricting to the diagonal group. Briefly speaking, conditions 1 and 2 ensure that we first restrict to monomial matrices. Conditions 3 and 4 prevent non-trivial permutations due to the following. As we assume n is larger than some constant, by Observation 5.9, $\sigma \in S_n$ either fixes 6 elements in $[n]$, or moves a set of 6 elements to another, disjoint, set of 6 elements. Condition 3 ensures that the second case could not happen. Condition 4 ensures that in the first case, the only possible invertible matrices that “preserves” the matrices G_i for $i \in P$ when multiplying from the left are scalar matrices.

We now prove Proposition 5.6, and this requires the following observation.

Observation 5.9. *Let $n \geq 23$. Then any permutation $\sigma \in S_n$ either fixes a set of 6 points $P \subseteq [n]$, or moves a set of 6 points $P \subseteq [n]$ to another set of 6 points $Q \subseteq [n]$ such that these two sets are disjoint.*

PROOF. Suppose σ fixes at most 5 points. Then there are at least 18 points that are not fixed by σ . Suppose σ has t non-trivial cycles of length $l_1, \dots, l_t$, such that $\sum_i l_i \geq 18$. For a cycle $(p_1, \dots, p_s)$, we can choose those points with odd indices, namely $p_1, p_3, \dots, p_{2\lfloor s/2 \rfloor - 1}$ and put them in P , and those points with even indices, namely $p_2, p_4, \dots, p_{2\lfloor s/2 \rfloor}$ in Q . Do this for every cycle, we obtain the desired P and Q . The worst case is when every cycle is of length 3. Since there are at least 18 points not fixed by σ , P is of size ≥ 6 . $\square$

PROOF OF PROPOSITION 5.6. Recall that we construct such $\tilde{A}$ and $\tilde{B}$ from A and B , respectively, using the method in Section 5.1. Let $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{B}}$ be alternating matrix spaces in $\Lambda(7n, q)$, spanned by the frontal slices of $\tilde{A}$ and $\tilde{B}$, respectively.

We want to show that $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{B}}$ are isometric if and only if $\mathcal{A}$ and $\mathcal{B}$ are isometric via diagonal matrices. The if direction is straightforward. Suppose there exist $P = \text{diag}(\alpha_1, \dots, \alpha_n) \in \text{diag}(n, q)$ and $Q \in \text{GL}(m, q)$ such that $P^t A P = B^Q$. Let $\tilde{P} = \begin{bmatrix} P & 0 \\ 0 & I_{6n} \end{bmatrix} \in \text{GL}(7n, q)$. Let $\tilde{Q} = \begin{bmatrix} Q & 0 \\ 0 & Q' \end{bmatrix} \in \text{GL}(m + 4n^2, q)$, where $Q' = \text{diag}(\alpha_1 I_{4n}, \dots, \alpha_n I_{4n})$.

Then it is easy to verify that $\tilde{P}^t \tilde{A} \tilde{P} = \tilde{B} \tilde{Q}$.

Now we turn to the only if direction. If $\tilde{A}$ and $\tilde{B}$ are isometric, then there exists $\tilde{P} \in \text{GL}(7n, q)$ and $\tilde{Q} \in \text{GL}(m + 4n^2, q)$, such that $\tilde{P}^t \tilde{A} \tilde{P} = \tilde{B} \tilde{Q}$. Let $\tilde{P} = \begin{bmatrix} P_{1,1} & P_{1,2} \\ P_{2,1} & P_{2,2} \end{bmatrix}$, where $P_{1,1}$ is of size $n \times n$. It can be checked easily, from the lateral viewpoint, that $P_{1,2} = 0$. As if not, then some H_i would appear in one of the last $6n$ lateral slices in $\tilde{A} \tilde{P}$. This would set this slice to be of rank $\geq 4n$ by condition (1), which contradicts that the corresponding lateral slice of $\tilde{B} \tilde{Q}$ is of rank $\leq n$. It follows that $P_{1,1} \in \text{GL}(n, q)$ and $P_{2,2} \in \text{GL}(6n, q)$.

We first claim that $P_{1,1}$ has to be a monomial matrix. If not, suppose the $P_{1,1}(i, j)$ and $P_{1,1}(i, k)$ are non-zero, $j \neq k$. Then the i th lateral slice of $\tilde{A} \tilde{P}$ contains two distinct H_j and H_k as submatrices. By condition (2), this slice is of rank $\geq 6n$. On the other hand, each lateral slice of $\tilde{B} \tilde{Q}$ is of the same rank as $\tilde{B}$ (as $\tilde{Q}$ does not change the ranks of lateral slices), which by Fact 5.1 is $\leq 5n$. This is a contradiction, showing that $P_{1,1}$ must be a monomial matrix.

We further claim that $P_{1,1}$ has to be a diagonal matrix. If not, then suppose the non-trivial permutation underlying $P_{1,1}$ is $\sigma \in S_n$. Since we assumed n is larger than some constant, by Observation 5.9, one of the following two cases has to happen.

- $\exists \{i_1, \dots, i_6\} \subseteq [n], \{j_1, \dots, j_6\} \subseteq [n], |\{i_1, \dots, i_6\} \cup \{j_1, \dots, j_6\}| = 12$, such that $\sigma(i_k) = j_k$ for $k \in [6]$. We then claim the following.

Claim 5.10. *For $\tilde{P}^t \tilde{A} \tilde{P} = \tilde{B} \tilde{Q}$ to hold, a necessary condition is that $\forall k \in [6], P_{2,2} H_{j_k}$ and H_{i_k} have the same linear span.*

PROOF. To see this, note that the i_k th lateral slice of $\tilde{P}^t \tilde{A} \tilde{P}$ is the j_k th lateral slice of $\tilde{P}^t \tilde{A}$ (up to a scalar multiple). It is equal to the i_k th lateral slice of $\tilde{B} \tilde{Q}$. Then $\tilde{P}^t$ acts on the left on the j_k th lateral slice of $\tilde{A}$. Noting that $P^t = \begin{bmatrix} P_{1,1}^t & P_{2,1}^t \\ 0 & P_{2,2}^t \end{bmatrix}$ and the j_k th lateral slice of $\tilde{A}$ is $C'_{j_k} = \begin{bmatrix} C_{j_k} & 0 \\ 0 & G_{j_k} \end{bmatrix}$, we see that $P^t C'_{j_k} = \begin{bmatrix} * & * \\ 0 & P_{2,2}^t G_{j_k} \end{bmatrix}$. (Here, C_i and G_i are defined in Section 5.1.) On the other hand, we see that the i_k th lateral slice of $\tilde{B} \tilde{Q}$ is the i_k th lateral slice of $\tilde{B}$ multiplied from the right by $\tilde{Q}$. Our claim follows then by comparing the last $6n$ rows. $\square$

But the condition (3) excludes the existence of such $P_{2,2}$, so this cannot happen.

- $\exists \{i_1, \dots, i_6\} \subseteq [n], i_k$ all different, such that $\sigma(i_k) = i_k$. In this case, for $\tilde{P}^t \tilde{A} \tilde{P} = \tilde{B} \tilde{Q}$ to hold, by the same argument as in the proof of Claim 5.10, a necessary condition is that $P_{2,2} H_{i_k}$ and H_{i_k} have the same linear span. Then the condition (4) ensures that $P_{2,2} = \lambda I_{6n}$ for some $\lambda \neq 0 \in \mathbb{F}$ in this setting. Then because σ is non-trivial, σ moves some $i \in [n]$ to $j \in [n], i \neq j$. By comparing the j th lateral slice of $\tilde{P}^t \tilde{A}$ and the i th lateral slice of $\tilde{B} \tilde{Q}$, $P_{2,2} H_i = \lambda H_i$ and H_j have the same linear span, which is not possible because the condition (2) ensures that H_i and H_j span different subspaces.

We then have shown that $P_{1,1}$ must be a diagonal matrix. By comparing the top-left-front sub-tensors of size $n \times n \times m$ of $\tilde{P}^t \tilde{A} \tilde{P}$ and $\tilde{B} \tilde{Q}$, we arrive at the desired conclusion that $\mathcal{A}$ and $\mathcal{B}$ are isometric via the diagonal matrix $P_{1,1}$. $\square$

Remark 5.11. If we only attach the diagonal restriction gadget to the first k slices (see Remark 5.2), then the above proof can be adapted to show that: $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{B}}$ are isometric, if and only if, $\mathcal{A}$ and $\mathcal{B}$ are isometric via $P = \begin{bmatrix} D & 0 \\ E & F \end{bmatrix}$ where D is a $k \times k$ diagonal matrix.

5.2.2 Using the gadget for counting-to-decision reduction. The strategy follows closely the counting to decision reduction for graph isomorphism.

We first review the strategy for counting to decision reduction for graph isomorphism [56]. Suppose we are given two graphs with the vertex set being $[n]$, i.e. $G, H \subseteq \binom{[n]}{2}$. We first use the decision oracle to decide whether G and H are isomorphic. If not, the number of isomorphisms is 0. If so, we turn to compute the order of $\text{Aut}(G)$. Let $A = \text{Aut}(G)$. For $i \in [n]$, let $A_i = \{\sigma \in A : \forall 1 \leq j \leq i, \sigma(j) = j\}$. Set $A_0 = A$. We then have the tower of subgroups $A_0 \geq A_1 \geq \dots \geq A_n = \{\text{id}\}$. The order of A_0 is then the product of $[A_i : A_{i+1}]$, the index of A_{i+1} in A_i , for $i = 0, 1, \dots, n-1$. Let G_i be the graph with the first i vertices in G individualized. Then $\text{Aut}(G_i) \cong A_i$. To compute $[A_i : A_{i+1}]$, we note that it is equal to the size of the orbit of the vertex $i+1$ under A_i . For each $j \geq i+1$, construct from G_i two graphs G'_i and G''_i as follows. In G'_i , individualize $i+1$, and in G''_i , individualize j . Then j is in the orbit of $i+1$ under A_i if and only if G'_i and G''_i are isomorphic. Enumerating over $j \geq i+1$ gives us the size of the orbit of $i+1$ under A_i . This finishes an overview of the idea for counting to decision reduction for graph isomorphism.

We then apply the above strategy to get a counting to decision reduction for alternating matrix space isometry to prove Proposition 5.7.

PROOF OF PROPOSITION 5.7. Our goal is to compute the number of isometries from $\mathcal{A}$ to $\mathcal{B}$, where $\mathcal{A}, \mathcal{B} \leq \Lambda(n, q)$ are of dimension m . First, we use the decision oracle first to decide whether $\mathcal{A}$ and $\mathcal{B}$ are isometric. If not, the number of isometries is 0. If so, we need to calculate the order of the autometry group of $\mathcal{A}$, $\text{Aut}(\mathcal{A})$, that is, the set of self-isometries $\mathcal{A} \rightarrow \mathcal{A}$ as a subgroup of $\text{GL}(n, q)$. To do that, we first randomly sample n $6n \times 4n$ matrices $H_1, \dots, H_n$ over $\mathbb{F}_q$, and verify whether they form a rigid matrix tuple using Remark 5.4. Note that this is where the algorithm needs to be a Las Vegas algorithm.

Let $A = \text{Aut}(\mathcal{A})$. Recall that e_i denotes the i th standard basis vector in $\mathbb{F}_q^n$. For $i \in [n]$, let $A_i = \{T \in A : \forall 1 \leq j \leq i, T(e_j) = \lambda_j e_j, \lambda_j \neq 0 \in \mathbb{F}_q\}$. Note that $A_n = A \cap \text{diag}(n, q)$. We can calculate the order of A_n in time $q^{O(n)}$ by brute-force, i.e., enumerating all invertible diagonal matrices. Set $A_0 = A$. We then have the tower of subgroups $A_0 \geq A_1 \geq \dots \geq A_n$.

To compute the order of A_0 , it is enough to compute $[A_i : A_{i+1}]$. Note that for $T, T' \in A_i$, $TA_{i+1} = T'A_{i+1}$ as left cosets in A_i if and only if $T(e_{i+1}) = \lambda T'(e_{i+1})$ for some $\lambda \neq 0 \in \mathbb{F}_q$. So $[A_i : A_{i+1}]$ is equal to the size of the orbit of e_{i+1} under A_i in the projective space. Let $v \in \mathbb{F}_q^n$. To test whether v is in the orbit of e_{i+1} under A_i in the projective space, we transform $\mathcal{A}$ by $P^t \cdot P$, where $P \in \text{GL}(n, q)$ sends e_{i+1} to v and e_j to e_j for $j \neq i+1$, to get $\mathcal{A}'$. We then add the diagonal restriction gadget to the first $i+1$ lateral slices and the first $i+1$ horizontal slices of $\mathcal{A}$ and $\mathcal{A}'$ (see Remark 5.2), to obtain $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{A}}'$ respectively. Then feed $\mathcal{A}$ and $\mathcal{A}'$ to the decision oracle. By the functionality of the diagonal restriction gadget (Proposition 5.6 and Remark 5.11), v is in the orbit of e_{i+1} in the projective space if and only if $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{A}}'$ are isometric. Enumerating $v \in \mathbb{F}_q^n$ up to scalar multiples gives us the size of the orbit of e_{i+1} under A_i in the projective space. This finishes the description of the algorithm.

A small caveat in the above is that our gadget requires n is larger than some constant, so we cannot start from A_0 at the beginning. This issue can be resolved by noting that the order of A_c , for any constant c , can be computed in time $q^{O(n)}$, by enumerating all possible images of $e_1, \dots, e_c$ in time $q^{O(n)}$, adding the diagonal restriction gadget, and utilizing the decision oracle. $\square$

5.2.3 Random H_i 's satisfy the requirements when $q = n^{\Omega(1)}$. We now prove Proposition 5.5, and for this we need the following facts.

Fact 5.12. (1) Given $a_i \in \mathbb{R}$, $0 \leq a_i \leq 1$, $i \in [m]$, $\prod_{i \in [m]} (1 - a_i) \geq 1 - \sum_{i \in [m]} a_i$.

(2) Let $m, N \in \mathbb{N}$ and $1 \leq m \leq N$. A random matrix $A \in M(N \times m, q)$ is of rank m with probability $\geq 1 - 2/q^{N-m+1}$.

(3) For $d \leq \mathbb{N}$, $0 \leq d \leq n$, the number of dimension- d subspaces of $\mathbb{F}_q^n$ is equal to the Gaussian binomial coefficient

$$\binom{n}{d}_q := \frac{(q^n - 1) \cdot (q^n - q) \cdot \dots \cdot (q^n - q^{d-1})}{(q^d - 1) \cdot (q^d - q) \cdot \dots \cdot (q^d - q^{d-1})}.$$

(4) The Gaussian binomial coefficient satisfies:

$$q^{(n-d)d} \leq \binom{n}{d}_q \leq q^{(n-d)d+d}.$$

(5) For $d \in \mathbb{N}$, the number of complement subspaces of a fixed dimension- d subspace of $\mathbb{F}_q^n$ is $q^{d(n-d)}$.

PROOF. (1) is clear. For (2), $\Pr[\text{rk}(A) = m] = (1 - \frac{1}{q^N}) \cdot (1 - \frac{q}{q^N}) \cdot \dots \cdot (1 - \frac{q^{m-1}}{q^N})$. By (1), we have $\Pr[\text{rk}(A) = m] \geq 1 - \sum_{i=N-m+1}^N \frac{1}{q^i} = 1 - \frac{1}{q^{N-m+1}} - \sum_{i=N-m+2}^N \frac{1}{q^i} \geq 1 - \frac{2}{q^{N-m+1}}$. (3) is classical; see e.g. [22]. For (4), it is because $q^{n-d} \leq \frac{q^n - q^i}{q^d - q^i} \leq q^{n-d+1}$. (5) is not hard to derive; see e.g. [23]. $\square$

In the following we will encounter random matrices over $\mathbb{F}_q$ as well as random subspaces in $\mathbb{F}_q^n$. There is a subtle point which we want to clarify now. Let $m \leq n$. Note that there are $\binom{n}{m}_q$ subspaces of $\mathbb{F}_q^n$ of dimension m , and there are $N_1 = (q^n - 1) \cdot \dots \cdot (q^n - q^{m-1})$ rank- m matrices of size $n \times m$. It can be seen easily that each m -dimensional subspace V of $\mathbb{F}_q^n$ has $N_2 = (q^m - 1) \cdot \dots \cdot (q^m - q^{m-1})$ many representations as rank- m matrices of size $n \times m$, i.e. the columns of the matrix span V . It follows that we can work with random rank- m matrices of size $n \times m$ as if we are working with random m -dimensional subspaces of $\mathbb{F}_q^n$. Such correspondences will be used implicitly for other structures, including direct sum decompositions.

Now let us get back to our question. We shall show that a random choice of $H_i, i \in [n]$, would form a rigid tuple. We will prove that for conditions $k = 1, 2, 3$,

$$\Pr[\text{random } H_i \text{ not satisfy condition } k] \leq \frac{n^{O(1)}}{q^{\Omega(n)}}.$$

Once these hold, by a union bound, we have

$$\Pr[\exists i \in [3], \text{random } H_i \text{ not satisfy condition } i] \leq \frac{n^{O(1)}}{q^{\Omega(n)}}.$$

For condition (4), we will prove that

$$\Pr[\text{random } H_i \text{ not satisfy condition 4} \mid H_i \text{ satisfy conditions 1, 2, 3}] \leq \frac{n^{O(1)}}{q^{\Omega(1)}}.$$

This then would allow us to conclude that when $q = n^{\Omega(1)}$, random H_i 's form a rigid matrix tuple.

We examine the first three conditions one by one.

- (1) For condition (1), by Fact 5.12 (2), we have $\Pr[\exists i \in [n], \text{rk}(H_i) < 4n] \leq n \cdot \Pr[\text{rk}(H_i) < 4n] \leq \frac{2n}{q^{2n+1}}$.
- (2) For condition (2), noting that the block matrix $(H_i H_j)$ is a random $6n \times 8n$ matrix over $\mathbb{F}_q$, by Fact 5.12 (2), we have $\Pr[\exists i \neq j \in [n], \text{rk}((H_i H_j)) < 6n] \leq \binom{n}{2} \cdot \frac{2}{q^{8n-6n+1}} \leq \frac{n^2}{q^{2n+1}}$.
- (3) For condition (3), let $I = (H_{i_1} \dots H_{i_6})$, and $J = (H_{j_1} \dots H_{j_6})$. We see that C is non-empty if and only if there exists $L \in \text{GL}(6n, q)$ and $R_k \in \text{GL}(4n, q)$, $k \in [6]$, such that $LH_{i_k}R_k = H_{j_k}$. Note that the orbit of I under this group action is of size at most $q^{(6n)^2+6 \cdot (4n)^2} = q^{132n^2}$. Since i_k and j_ℓ are all different, the probability of J belonging to this orbit is $\leq \frac{q^{132n^2}}{q^{144n^2}} = \frac{1}{q^{12n^2}}$. We then have $\Pr[\exists i_k, j_k \in [n], k \in [6], i_k, j_k \text{ all different}, C \neq \emptyset] \leq \binom{n}{12} \frac{2}{q^{12n^2}} \leq \frac{n^{12}}{q^{12n^2}}$.

We now focus on condition (4). For condition (4), we first assume that the conditions (1) and (2) as above hold. Then V_i 's are random $4n$ -dimensional subspaces of $\mathbb{F}_q^{6n}$. Note that

$$\Pr[\exists i_k \in [n], k \in [6], i_k \text{ all different}, S \text{ non-scalar}] \leq n^6 \cdot \Pr[S \text{ non-scalar stabilizer for } V_1, \dots, V_6].$$

So we turn to study $\Pr[S \text{ non-scalar stabilizer for } V_1, \dots, V_6]$, and will show that it is $\leq \frac{1}{q^{\Omega(1)}}$.

Let $U_1 = V_1 \cap V_2$, $U_2 = V_2 \cap V_3$, and $U_3 = V_1 \cap V_3$. Let $W_1 = V_4 \cap V_5$, $W_2 = V_5 \cap V_6$, and $W_3 = V_4 \cap V_6$. Since conditions (1) and (2) hold, we have $\dim(U_i) = \dim(W_i) = 2n$. We claim that with probability $\geq 1 - 2/q$, $\mathbb{F}_q^{6n} = U_1 \oplus U_2 \oplus U_3$, i.e., $U_1 \cup U_2 \cup U_3$ span $\mathbb{F}_q^{6n}$. This can be seen as follows. Since we assumed conditions (1) and (2), this happens if

and only if $V_1 \cap V_2$ and V_3 together span $\mathbb{F}_q^{6n}$. Therefore we calculate, using Fact 5.12 (1), (3), and (5), that

$$\begin{aligned}
& \Pr[V_3 \text{ is a complement subspace of } V_1 \cap V_2] \\
&= q^{2n \cdot 4n} / \binom{6n}{4n}_q = \frac{(q^{6n} - q^{2n})(q^{6n} - q^{2n+1}) \dots (q^{6n} - q^{6n-1})}{(q^{6n} - 1)(q^{6n} - q) \dots (q^{6n} - q^{4n-1})} \\
&\geq \frac{(q^{6n} - q^{2n})(q^{6n} - q^{2n+1}) \dots (q^{6n} - q^{6n-1})}{q^{6n} \cdot q^{6n} \dots q^{6n}} = (1 - 1/q^{4n})(1 - 1/q^{4n-1}) \dots (1 - 1/q) \\
&\geq 1 - \sum_{i=1}^{4n} 1/q^i \geq 1 - 2/q.
\end{aligned}$$

It follows that with probability $\geq 1 - 4/q$, we can assume in addition that W_i form a direct sum decomposition of $\mathbb{F}_q^{6n}$.

Therefore, we turn to bound the probability that there exists a non-scalar invertible matrix stabilizing these two direct sum decompositions of $\mathbb{F}_q^{6n}$. By showing that, under suitable conditions, this probability is at most $1/n^{\Omega(1)}$, we conclude that a random choice of subspaces works as our gadget with probability $1 - 1/n^{\Omega(1)}$, which suffices for a Las Vegas algorithm. Since i_k are all different, the two direct sum decompositions $U_1 \oplus U_2 \oplus U_3$ and $W_1 \oplus W_2 \oplus W_3$ are independent. So we can assume that U_i is spanned by those standard basis vectors $\vec{e}_{2n(i-1)+1}, \dots, \vec{e}_{2ni}$, $i = 1, 2, 3$.

The group that stabilizes this direct sum decomposition $U_1 \oplus U_2 \oplus U_3$ consists of $\begin{bmatrix} D_1 & 0 & 0 \\ 0 & D_2 & 0 \\ 0 & 0 & D_3 \end{bmatrix} \in \text{GL}(6n, \mathbb{F}_q)$

where D_i is of size $2n \times 2n$.

The question then becomes to bound the probability for a random $W_1 \oplus W_2 \oplus W_3$ to be stabilized by a non-scalar matrix of the above form. This can be formulated as the following linear algebraic problem. (Recall the correspondence between random m -dimensional subspaces and random rank- m matrices as discussed at the

beginning of the subsection.) Let $W = \begin{bmatrix} W_{11} & W_{12} & W_{13} \\ W_{21} & W_{22} & W_{23} \\ W_{31} & W_{32} & W_{33} \end{bmatrix} \in \text{GL}(6n, q)$ be a block matrix where W_{ij} is of size

$2n \times 2n$. Suppose the columns of $\begin{bmatrix} W_{1i} \\ W_{2i} \\ W_{3i} \end{bmatrix}$ span W_i . Then $D = \text{diag}(D_1, D_2, D_3)$ stabilizes $W_1 \oplus W_2 \oplus W_3$ if and only if there exists a block diagonal matrix $E = \text{diag}(E_1, E_2, E_3)$, $E_i \in \text{GL}(2n, q)$, such that

$$\begin{bmatrix} D_1 & 0 & 0 \\ 0 & D_2 & 0 \\ 0 & 0 & D_3 \end{bmatrix} \begin{bmatrix} W_{11} & W_{12} & W_{13} \\ W_{21} & W_{22} & W_{23} \\ W_{31} & W_{32} & W_{33} \end{bmatrix} = \begin{bmatrix} W_{11} & W_{12} & W_{13} \\ W_{21} & W_{22} & W_{23} \\ W_{31} & W_{32} & W_{33} \end{bmatrix} \begin{bmatrix} E_1 & 0 & 0 \\ 0 & E_2 & 0 \\ 0 & 0 & E_3 \end{bmatrix}. \quad (8)$$

Note that each direct sum decomposition $W_1 \oplus W_2 \oplus W_3$, $\dim(W_i) = 2n$, has $6 \cdot |\text{GL}(2n, q)|^3$ such matrix representations. (The factor 6 takes care of the orders of the three summands.) So the question becomes to bound the probability for a random invertible matrix to have a non-scalar D and E satisfying Equation 8.

First, note that Equation 8 holds if and only if $D_i W_{i,j} = W_{i,j} E_j$ for $i, j \in [3]$.

Claim 5.13. *When $q = \Omega(1)$, we have $\Pr[\forall i, j \in [3], \text{rk}(W_{i,j}) = 2n] \geq 1 - \frac{20}{q}$.*

PROOF. Let us work in the setting when W is a random matrix, not necessarily invertible. Then $\Pr[\text{rk}(W) = 6n] \geq 1 - \frac{2}{q}$. For any $i, j \in [3]$, $\Pr[\text{rk}(W_{i,j}) < 2n] \leq \frac{2}{q}$, so $\Pr[\exists i, j \in [3], \text{rk}(W_{i,j}) < 2n] \leq \frac{18}{q}$. It follows that $\Pr[\exists i, j \in [3], \text{rk}(W_{i,j}) < 2n \mid \text{rk}(W) = 6n] = \Pr[\exists i, j \in [3], \text{rk}(W_{i,j}) < 2n \wedge \text{rk}(W) = 6n] / \Pr[\text{rk}(W) = 6n] \leq \frac{18/q}{1-2/q} = \frac{18}{q-2} \leq \frac{20}{q}$, where the last inequality uses that $q = \Omega(1)$. $\square$

So we assume that $\text{rk}(W_{i,j}) = 2n$ for all $i, j \in [3]$ in the following, with a loss of probability $\leq \frac{20}{q}$.

For $i \in [3]$, by $D_i W_{ii} = W_{ii} E_i$, we have $D_i = W_{ii} E_i W_{ii}^{-1}$. For $i \neq j$, by $(W_{jj} E_j W_{jj}^{-1}) W_{ji} = D_j W_{ji} = W_{ji} E_i$, we have $E_j = W_{jj}^{-1} W_{ji} E_i W_{ji}^{-1} W_{jj}$. Again for $i \neq j$, we have $W_{ii} E_i W_{ii}^{-1} W_{ij} = D_i W_{ij} = W_{ij} E_j = W_{ij} W_{jj}^{-1} W_{ji} E_i W_{ji}^{-1} W_{jj}$. It follows that

$$\forall i, j \in [3], i \neq j, E_i W_{ii}^{-1} W_{ij} W_{jj}^{-1} W_{ji} = W_{ii}^{-1} W_{ij} W_{jj}^{-1} W_{ji} E_i.$$

In particular, E_3 commutes with $X = W_{33}^{-1} W_{32} W_{22}^{-1} W_{23}$ and $Y = W_{33}^{-1} W_{31} W_{11}^{-1} W_{13}$. Since W_{ij} are independent random invertible matrices, X and Y are independent random invertible matrices. We now resort to the following classical result.

Theorem 5.14 ([44], cf. also [43, Theorem 3.3] and [26, The paragraph after Theorem 1.1]). *Let X and Y be two random matrices in $\text{GL}(n, q)$. Then the probability of X and Y not generating a group containing $\text{SL}(n, q)$ is $\leq \frac{1}{q^{\Omega(n)}}$.*

It follows that E_3 belongs to the centralizer of G , so E_3 must be a scalar matrix. Then note that D_i 's and other E_i 's are all conjugates of E_3 . So we have $\forall i \in [3], D_i = E_i = \lambda I_{2n}$ for some $\lambda \neq 0 \in \mathbb{F}_q$.

Summarizing the above, we have

$$\begin{aligned} & \Pr[S \text{ non-scalar for } V_1, \dots, V_6] \\ & \leq \Pr[S \text{ non-scalar for } V_i \wedge \mathbb{F}_q^{6n} = U_1 \oplus U_2 \oplus U_3 = W_1 \oplus W_2 \oplus W_3] + \frac{4}{q} \\ & \leq \Pr[S \text{ non-scalar for } V_i \mid \mathbb{F}_q^{6n} = U_1 \oplus U_2 \oplus U_3 = W_1 \oplus W_2 \oplus W_3] + \frac{4}{q} \\ & \leq \Pr[D \text{ non-scalar for } W \wedge \forall i, j \in [3], \text{rk}(W_{ij}) = 2n] + \frac{20}{q} + \frac{4}{q} \\ & \leq \Pr[D \text{ non-scalar for } W \mid \forall i, j \in [3], \text{rk}(W_{ij}) = 2n] + \frac{24}{q} \\ & \leq \frac{1}{q^{\Omega(n)}} + \frac{24}{q} \\ & \leq \frac{1}{q^{\Omega(1)}}. \end{aligned}$$

This concludes the proof of Proposition 5.5. □

6 APPLICATION TO p -GROUP ISOMORPHISM, USING CONSTRUCTIVE BAER AND LAZARD CORRESPONDENCES

The applications to p -GROUP ISOMORPHISM rely on the following well-known connections between alternating bilinear maps and Lie algebras on the one hand, and p -groups of “small” class on the other. We present these connections here, partly for audiences not from computational group theory, and partly because we will need to address some computational aspects of these procedures. We begin with some preliminaries.

6.1 Preliminaries

TI-completeness. As the proof of Thm. P in Section 6.3.1 uses a result on TI-completeness from [35], here we recall the definition of TI; see Definition 3.1 for the d -TENSOR ISOMORPHISM problem.

Definition 6.1 ($d\text{TI}$, TI). For any field $\mathbb{F}$, $d\text{TI}_{\mathbb{F}}$ denotes the class of problems that are polynomial-time Turing (Cook) reducible to d -TENSOR ISOMORPHISM over $\mathbb{F}$. Also let $\text{TI}_{\mathbb{F}} = \bigcup_{d \geq 1} d\text{TI}_{\mathbb{F}}$.

The relationship between TI over different fields remains an intriguing open question [35], but here we will only need TI over $\mathbb{F}_p$. One of the main results of [35] is that $\text{TI} = d\text{TI}$ for any fixed $d \geq 3$.

Algebras and their algorithmic representations. A Lie algebra $\mathcal{A}$ consists of a vector space V and a bilinear map $[\cdot, \cdot] : V \times V \rightarrow V$ that is alternating ($[v, v] = 0$ for all $v \in V$; this is equivalent to skew-symmetry $[u, v] = -[v, u]$ in characteristic not 2) and satisfies the Jacobi identity $[x, [y, z]] + [z, [x, y]] + [y, [z, x]] = 0$. The Jacobi identity is essentially the “derivative” of associativity.

After choosing an ordered basis $(b_1, \dots, b_n)$ where $b_i \in \mathbb{F}^n$ of $V \cong \mathbb{F}^n$, this bilinear map $[\cdot, \cdot]$ can be represented by an $n \times n \times n$ 3-way array A , such that $[b_i, b_j] = \sum_{k \in [n]} A(i, j, k) b_k$. This is the structure constant representation of $\mathcal{A}$. Algorithms for Lie algebras have been studied intensively in this model, e. g., [24, 41].

It is also natural to consider matrix spaces that are closed under commutator. More specifically, let $\mathcal{A} \leq M(n, \mathbb{F})$ be a matrix space. If $\mathcal{A}$ is closed under commutator, that is, for any $A, B \in \mathcal{A}$, $[A, B] = AB - BA \in \mathcal{A}$, then $\mathcal{A}$ is a matrix Lie algebra with the product being the commutator. (Protip: one way to remember the Jacobi identity is to derive it as the natural identity among nested commutators of three matrices.) Algorithms for matrix Lie algebras have also been studied, e. g., [27, 39, 41].

6.2 Constructive Baer Correspondence and Theorems A and B

Let us review Baer’s correspondence [8], which connects alternating bilinear maps with p -groups of class 2 and exponent p . Let P be a p -group of class 2 and exponent p , $p > 2$. Suppose the commutator subgroup $[P, P] \cong \mathbb{Z}_p^m$ and $P/[P, P] \cong \mathbb{Z}_p^n$. Then the commutator map $[\cdot, \cdot] : P/[P, P] \times P/[P, P] \rightarrow [P, P]$ is an alternating bilinear map. Conversely, let $\phi : \mathbb{Z}_p^n \times \mathbb{Z}_p^n \rightarrow \mathbb{Z}_p^m$ be an alternating bilinear map. Then a p -group of class 2 and exponent p , denoted as P_ϕ can be defined as follows. The group elements are from $\mathbb{Z}_p^n \times \mathbb{Z}_p^m$, and the group product $\cdot$ is defined as

$$(u, v) \cdot (u', v') = (u + u', v + v' + \frac{1}{2}\phi(u, u')).$$

We say that $(A, B) \in \text{GL}(n, p) \times \text{GL}(m, p)$ is a pseudo-autometry of ϕ , if $\phi(u, v) = B\phi(Au, Av)$ for all $u, v \in \mathbb{Z}_p^n$. Clearly, there is a one-to-one correspondence between automorphisms of P_ϕ and pseudo-autometries of ϕ .

We then state a lemma which can be viewed as a constructive version of Baer’s correspondence, communicated to us by James B. Wilson.

Lemma 6.2 (Constructive version of Baer’s correspondence for matrix groups). *Let p be an odd prime. Over the finite field $\mathbb{F} = \mathbb{F}_{p^e}$, ALTERNATING MATRIX SPACE ISOMETRY is equivalent to GROUP ISOMORPHISM for matrix groups over $\mathbb{F}$ that are p -groups of class 2 and exponent p . More precisely, there are functions computable in time $\text{poly}(n, m, \log |\mathbb{F}|)$:*

- $G : \Lambda(n, \mathbb{F})^m \rightarrow M(n + m + 1, \mathbb{F})^{n+m}$ and
- $\text{Alt} : M(n, \mathbb{F})^m \rightarrow \Lambda(m, \mathbb{F})^{O(m^2)}$

such that: (1) for an alternating bilinear map A , the group generated by $G(A)$ is the Baer group corresponding to A , (2) G and Alt are mutually inverse, in the sense that the group generated by $G(\text{Alt}(M_1, \dots, M_m))$ is isomorphic to the group generated by $M_1, \dots, M_m$, and conversely $\text{Alt}(G(A))$ is pseudo-isometric to A .

PROOF. First, let G be a p -group of class 2 and exponent p given by m generating matrices of size $n \times n$ over $\mathbb{F}$. Then from the generating matrices of G , we first compute a generating set of $[G, G]$, by just computing all the commutators of the given generators. We can then remove those redundant elements from this generating set in time $\text{poly}(\log |[G, G]|, \log |\mathbb{F}|)$, using Luks’ result on computing with solvable matrix groups[53]. We then compute a set of representatives of a non-redundant generating set of $G/[G, G]$, again using Luks’s aforementioned result. From these data we can compute an alternating bilinear map representing the commutator map of G in time $\text{poly}(n, m, \log |\mathbb{F}|)$.

Conversely, let an alternating bilinear map be given by $\mathbf{A} = (A_1, \dots, A_m) \in \Lambda(n, \mathbb{F})^m$. From $\mathbf{A}$, for $i \in [n]$, construct $B_i = [A_1 e_i, \dots, A_m e_i] \in M(n \times m, \mathbb{F})$, where e_i is the i th standard basis vector of $\mathbb{F}^n$. That is, the j th column of B_i is the i th column of A_j . Then for $i \in [n]$, construct

$$\tilde{B}_i = \begin{bmatrix} 1 & e_i^t & 0 \\ 0 & I_n & B_i \\ 0 & 0 & I_m \end{bmatrix} \in \text{GL}(1 + n + m, \mathbb{F}),$$

where $e_i \in \mathbb{F}^n$, and for $j \in [m]$, construct

$$\tilde{C}_j = \begin{bmatrix} 1 & 0 & e_j^t \\ 0 & I_n & 0 \\ 0 & 0 & I_m \end{bmatrix} \in \text{GL}(1 + n + m, \mathbb{F}),$$

where $e_j \in \mathbb{F}^m$. Let $G(\mathbf{A})$ be the tuple consisting of the $\tilde{B}_i$ and the $\tilde{C}_j$, and let Γ be the group they generate. Then it can be verified easily that, Γ is isomorphic to the Baer group corresponding to the alternating bilinear map defined by $\mathbf{A}$. In particular, $[\Gamma, \Gamma] \cong \mathbb{F}^m \cong \mathbb{Z}_p^{em}$ (isomorphism of abelian groups), and $\Gamma/[\Gamma, \Gamma] \cong \mathbb{F}^n \cong \mathbb{Z}_p^{en}$. This construction can be done in time $\text{poly}(n, m, \log |\mathbb{F}|)$. $\square$

Given the above lemma, we can present search- and counting-to-decision reductions for testing isomorphism of a class of p -groups, proving Theorems A and B.

PROOF OF THEOREM A. The search-to-decision reduction follows from Theorem A', using the $q^{O(n+m)}$ -time algorithm, with the constructive version of Baer's Correspondence in the model of matrix groups over finite fields (Lemma 6.2).

In more detail, given Lemma 6.2 we can follow the procedure in the proof of Theorem A'. For the given p -groups, we compute their commutator maps. Then whenever we need to feed the decision oracle, we transform from the alternating bilinear map to a generating set of a p -group of class 2 and exponent p with this bilinear map as the commutator map. After getting the desired pseudo-isometry for the alternating bilinear maps, we can easily recover an isomorphism between the originally given p -groups. $\square$

PROOF OF THEOREM B. For the counting-to-decision reduction, we basically follow the above routine, but with a twist, because of the minor distinction between alternating matrix space isometry, and alternating bilinear map pseudo-isometry. Let us briefly explain this issue. Suppose from an alternating bilinear map $\phi : \mathbb{Z}_p^n \times \mathbb{Z}_p^n \rightarrow \mathbb{Z}_p^m$ we constructed a p -group of class 2 and exponent p P_ϕ , and there is a k -to-one correspondence between automorphisms of P_ϕ and pseudo-autometries of ϕ (we explain the value of k below). Let $(C_1, \dots, C_m) \in \Lambda(n, p)$ be a matrix representation of ϕ . If C_i 's are linearly independent, then for a pseudo-autometry $(A, B) \in \text{GL}(n, p) \times \text{GL}(m, p)$, given A there exists a unique B that makes (A, B) a pseudo-autometry. If C_i 's are not linearly independent, say the linear span of C_i 's is of dimension m' , then the number of B such that (A, B) is a pseudo-autometry is $|\text{GL}(m - m', p)|$. The counting to decision reduction for ALTERNATING MATRIX SPACE ISOMETRY computes the number of $A \in \text{GL}(n, p)$ so that there exists some $B \in \text{GL}(m, p)$ such that (A, B) is a pseudo-autometry. So it needs to be multiplied by a factor of $|\text{GL}(m - m', p)|$.

Furthermore, there are automorphisms of P_ϕ that act trivially on both $Z(P_\phi)$ and $P_\phi/Z(P_\phi)$, and hence correspond to the trivial pseudo-autometry of ϕ . Such automorphisms are in bijective correspondence with $\text{Hom}(\mathbb{Z}_p^n, \mathbb{Z}_p^m)$, hence there are precisely p^{nm} of them—this is the factor of k mentioned above. For similar reasons, if the C_i span a space of dimension m' , we multiply by another factor of $p^{m'(m-m')}$ to get the number of automorphisms of P_ϕ . $\square$

6.3 Constructive Lazard's correspondence and Thm. P

The Lazard correspondence [49] is a correspondence between certain classes of groups and Lie algebras, which extends the usual correspondence between Lie groups and Lie algebras (say, over $\mathbb{R}$) to some groups and Lie algebras in positive characteristic. Here we state just enough to give a sense of it; for further details and exposition we refer to Khukhro's book [46] and Naik's thesis [62]. While Naik's thesis is quite long, it also includes a reader's guide, and collects many results scattered across the literature or well-known to the experts in one place, building the theory from the ground up and with many examples.

Recall that a *Lie ring* is an abelian group L equipped with a bilinear map $[\cdot, \cdot]$, called the Lie bracket, which is (1) alternating ($[x, x] = 0$ for all $x \in L$) and (2) satisfies the Jacobi identity $[x, [y, z]] + [y, [z, x]] + [z, [x, y]] = 0$ for all $x, y, z \in L$ (in some sense the "derivative" of the associativity equation). Let $L^1 = L$, and $L^{i+1} = [L, L^i]$, which is the subgroup (of the underlying additive group) generated by all elements of the form $[x, y]$ for $x \in L, y \in L^i$. Then L is *nilpotent* if $L^{c+1} = 0$ for some finite c ; the smallest such c is the *nilpotency class*. (Lie algebras are just Lie rings over a field.)

The correspondence between Lie algebras and Lie groups over $\mathbb{R}$ uses the Baker–Campbell–Hausdorff (BCH) formula to convert between a Lie algebra and a Lie group, so we start there. For non-commuting matrices X, Y , $e^X e^Y \neq e^{X+Y}$ in general (where the matrix exponential here is defined using the power series for e^x). Rather, using commutators $[A, B] = AB - BA$, we have

$$\exp(X) \exp(Y) = \exp \left(X + Y + \frac{1}{2}[X, Y] + \frac{1}{12}([X, [X, Y]] - [Y, [X, Y]]) - \frac{1}{24}[Y, [X, [X, Y]]] + \cdots \right),$$

where the remaining terms are iterated commutators that all involve at least 5 X s and Y s, and successive terms involve more and more. The BCH formula is a function of X, Y , that is given by the infinite summation inside the exponential on the RHS of the preceding equation. Applying the exponential function to a Lie algebra in characteristic zero yields a Lie group. The BCH formula can be inverted, giving the correspondence in the other direction.

In a nilpotent Lie algebra, the BCH formula has only finitely many nonzero terms, so issues of convergence disappear and we may consider applying the correspondence over finite fields or rings; the only remaining obstacle is that the denominators appearing in the formula must be units in the ring. It turns out that the correspondence continues to work in characteristic p so long as one does not need to use the p -th term of the BCH formula (which includes division by p), and the latter is avoided whenever a nilpotent group has class strictly less than p , or even when all subgroups generated by at most 3 elements have class strictly less than p . While the correspondence does apply more generally, here we only state the version for finite groups. For any fixed nilpotency class c , computing the Lazard correspondence is efficient in theory; for how to compute it in practice when the groups are given by polycyclic presentations, see [21].

Let $\mathbf{Grp}_{p,n,c}$ denote the set of finite groups of order p^n and class c , and let $\mathbf{Lie}_{p,n,c}$ denote the set of Lie rings of order p^n and class c . We note that for nilpotency class 2, the Baer correspondence is the same as the Lazard correspondence.

Theorem 6.3 (Lazard Correspondence for finite groups [49], see, e.g., [46, Ch. 9 & 10] or [62, Ch. 6]). *For any prime p and any $1 \leq c < p$, there are functions $\mathbf{log}: \mathbf{Grp}_{p,n,c} \leftrightarrow \mathbf{Lie}_{p,n,c} : \mathbf{exp}$ such that (1) $\mathbf{log}$ and $\mathbf{exp}$ are inverses of one another, (2) two groups $G, H \in \mathbf{Grp}_{p,n,c}$ are isomorphic if and only if $\mathbf{log}(G)$ and $\mathbf{log}(H)$ are isomorphic, and (3) if G has exponent p , then the underlying abelian group of $\mathbf{log}(G)$ has exponent p . More strongly, $\mathbf{log}$ is an isomorphism of categories $\mathbf{Grp}_{p,n,c} \cong \mathbf{Lie}_{p,n,c}$.*

Part (3) can be found as a special case of [62, Lemma 6.1.2].

For p -groups given by $d \times d$ matrices over the finite field $\mathbb{F}_{p^e}$, we will need one additional fact about the correspondence, namely that it also results in a Lie algebra of $d \times d$ matrices. (Being able to bound the dimension

of the Lie algebra and work with it in a simple linear-algebraic way seems crucial for our reduction to work efficiently.) In fact, the BCH correspondence is *easier* to see for matrix groups using the matrix exponential and matrix logarithm; most of the work for BCH and Lazard is to get the correspondence to work even *without* the matrices. In some sense, this is thus the “original” setting of this correspondence. Though it is surely not new, we could not find a convenient reference for this fact about matrix groups over finite fields, so we state it formally here.

Proposition 6.4 (cf. [46, Exercise 10.6]). *Let $G \leq \text{GL}(d, \mathbb{F}_{p^e})$ be a finite p -subgroup of exponent p , consisting of $d \times d$ matrices over a finite field of characteristic p . Then $\log(G)$ (from the Lazard correspondence) can be realized as a finite Lie subalgebra of $de \times de$ matrices over $\mathbb{F}_p$. Given a generating set for G of m matrices, a generating set for $\log(G)$ can be constructed in $\text{poly}(d, n, e \log p)$ time.*

Khukhro [46] gives the characteristic zero analogue of this result (minus the straightforward complexity analysis) for the full group of upper unitriangular matrices as Exercise 10.6. One way to see Proposition 6.4 is to use the characteristic zero result, apply the fact that these isomorphism are in fact equivalence of categories (and thus hold for subgroups/subalgebras as well), and note that the same formulae in characteristic zero apply in characteristic p so long as one never needs to divide by p . We now sketch the argument.

PROOF SKETCH. First we use the standard embedding of $\text{GL}(d, \mathbb{F}_{p^e})$ into $\text{GL}(de, \mathbb{F}_p)$ (replace each element by an $e \times e$ block which is the left regular representation of $\mathbb{F}_{p^e}$ acting on itself as an e -dimensional $\mathbb{F}_p$ -vector space), to realize G as a subgroup of $\text{GL}(de, \mathbb{F}_p)$. G is conjugate in $\text{GL}(de, \mathbb{F}_p)$ to a group of upper unitriangular matrices (upper triangular with all 1s on the diagonal); this is a standard fact that can be seen in several ways, for example, by noting that the group U of all upper unitriangular matrices in $\text{GL}(de, \mathbb{F}_p)$ is a Sylow p -subgroup, and applying Sylow’s Theorem. (Note that we do not need to do this conjugation algorithmically, though it is possible to do so [30, 39, 66]; this is only for the proof.) Thus we may write every $g \in G$ as $1 + n$, where the sum here is the ordinary sum of matrices, 1 denotes the identity matrix, and n is strictly upper triangular. To see that we can truncate the Taylor series for logarithm before the p -th term (thus avoiding needing to divide by p), note that $(1 + n)^p = 1$ since G is exponent p . We have $(1 + n)^p = 1^p + \binom{p}{1}n + \binom{p}{2}n^2 + \dots + \binom{p}{p-1}n^{p-1} + n^p$. Since these are matrices over a field of characteristic p , and $p \mid \binom{p}{i}$ for all $1 \leq i \leq p-1$, all the intermediate terms vanish and we have that $(1 + n)^p = 1^p + n^p$. Thus $1 = (1 + n)^p = 1 + n^p$, so we get that $n^p = 0$. Thus, in the the Taylor series for the logarithm

$$\log(1 + n) = n - \frac{n^2}{2} + \frac{n^3}{3} - \dots$$

the last nonzero term is $n^{p-1}/(p-1)$, so we may use this Taylor series even over $\mathbb{F}_{p^e}$.

The main things to check are that the set $\log(G) := \{\log(1 + n) : 1 + n \in G\}$ is closed under scalar multiplication, matrix addition, and matrix commutator $[X, Y] = XY - YX$. Suppose g_1, g_2 are matrices in G , and write them as $g_i = 1 + n_i$ ($i = 1, 2$), as above. We recall that, because $n_i^p = 0$ from above, the power series for both $\log$ and $\exp$ work to compute the matrix logarithm and exponential over $\mathbb{F}_{p^e}$, respectively, and that the usual rules of logarithms are satisfied for a single matrix A : whenever $A \in M_{de}(\mathbb{F}_p)$ satisfies $A^p = 0$, we have $\log \exp A = A$, $\exp \log(1 + A) = 1 + A$, $\exp(nA) = (\exp A)^n$ for $n \in \mathbb{Z}$, and $\log((1 + A)^n) = n \log(1 + A)$.

- **Scalar multiplication:** For $\alpha \in \mathbb{F}_p$, we show that $n \log(1 + n_1)$ is in $\log(G)$. This is easy to show, as it follows directly from the rules of logarithms just mentioned: $\alpha \log(1 + n_1) = \log((1 + n_1)^\alpha)$ where on the right-hand side we treat α as an integer in the range $[0, p-1]$.
- **Addition:** Let $x_i = \log(1 + n_i)$ for $i = 1, 2$. We want to show that $x_1 + x_2$ is in $\log(G)$, or equivalently that $\exp(x_1 + x_2) \in G$. This follows from the first inverse BCH formula h_1 , which satisfies $\exp(\hat{x}_1 + \hat{x}_2) = h_1(\exp(\hat{x}_1), \exp(\hat{x}_2))$ for $\hat{x}_i$ in the free nilpotent-of-class- c $\mathbb{F}_{p^e}$ -Lie algebra, and then we may apply the homomorphism from the latter algebra to the subalgebra of $M_n(\mathbb{F}_{p^e})$ generated by the n_i to see that the

same formula works. (We note, because a reviewer asked, that here we do not need this entire subalgebra to be in $\{g - 1 : g \in G\}$; the use of that subalgebra is just convenient for talking about algebra homomorphisms in the proof. Rather, it suffices that the preceding equation holds for these particular elements n_i , which are by definition of the form $g_i - 1$ for some matrices $g_i \in G$.)

- Commutator: $[\log(1 + n_1), \log(1 + n_2)]$. A similar argument as in the previous case works, using the second inverse BCH formula h_2 , which satisfies $\exp([\hat{x}_1, \hat{x}_2]) = h_2(\exp(\hat{x}_1), \exp(\hat{x}_2))$.

Equivalently, we may note that the derivation of the inverse BCH formula in [46, 62] uses a free nilpotent associative algebra as an ambient setting in which both the group (or rather, n such that $1 + n$ is in the group) and the corresponding Lie algebra live; in our case, we may replace the ambient free nilpotent associative algebra with the algebra of $de \times de$ strictly upper-triangular matrices over $\mathbb{F}_p$, and all the derivations remain the same, *mutatis mutandis*. See, for example, [46, p. 105, “Another remark...”]. $\square$

6.3.1 Class reduction in p -group isomorphism testing. Proposition 6.4 now allows us to prove Thm. P.

PROOF OF THM. P. By the Lazard correspondence (reproduced as Theorem 6.3) two p -groups of exponent p and class $c < p$ are isomorphic if and only if their corresponding $\mathbb{F}_p$ -Lie algebras are. By Proposition 6.4, we can construct a generating set for the corresponding $\mathbb{F}_p$ -Lie algebra by applying the power series for logarithm to the generating matrices of G . This Lie algebra is thus a subalgebra of $ne \times ne$ matrices over $\mathbb{F}_p$, so we can generate a basis for the entire Lie algebra (using the linear-algebra version of breadth-first search; its dimension is $\leq (ne)^2$) and compute its structure constants in time polynomial in n , m , and $e \log p$. Then use [31] to reduce isomorphism of Lie algebras to 3-TENSOR ISOMORPHISM, and then use the fact that isomorphism of p -groups of exponent p and class 2 given by a matrix generating set over $\mathbb{F}_p$ is TI-complete [35] to reduce to the latter problem. $\square$

7 CONCLUSION

In this paper, we gave first-of-their-kind results around search-to-decision, counting-to-decision, and reductions to hard instances in the context of GROUP ISOMORPHISM. We focused on p -groups of class 2 (or more generally small class) and exponent p , as these are widely believed to be the hardest cases of GPI. They also have the closest connection with tensors.

We view this paper as the second in a planned series, focusing on isomorphism problems for tensors, groups, polynomials, and related structures. Although GRAPH ISOMORPHISM (GI) is perhaps the most well-studied isomorphism problem in computational complexity—even going back to Cook’s and Levin’s initial investigations into NP (see [2, Sec. 1])—it has long been considered to be solvable in practice [57, 58], and Babai’s recent quasi-polynomial-time breakthrough is one of the theoretical gems of the last several decades [4]. However, several isomorphism problems for tensors, groups, and polynomials seem to be much harder to solve, both in practice—they’ve been suggested as difficult enough to support cryptography [42, 63]—and in theory: the best known worst-case upper bounds are barely improved from brute force (e.g., [52, 68]). As these problems arise in a variety of areas, from multivariate cryptography and machine learning, to quantum information and computational algebra, getting a better understanding of their complexity is an important goal with many potential applications.

In the first paper in this series [35], we showed that numerous such isomorphism problems from many research areas are equivalent under polynomial-time reductions, creating bridges between different disciplines. The TENSOR ISOMORPHISM (TI) problem turns out to occupy a central position among these problems, leading us to define the complexity class TI, consisting of those problems polynomial-time reducible to the TENSOR ISOMORPHISM problem. The gadgets and TI-completeness result from that first paper in some cases opened the door, and in other cases are used as subroutines, in the main results of the current paper.

Finally, we list here some additional questions that we find interesting and approachable. One question is whether our tensor-based methods here can be extended or combined with other methods to get analogous results in wider classes of groups; for isomorphism algorithms, something along these lines was proposed by Brooksbank, Grochow, Li, Wilson, & Qiao [13], but there are many interesting open questions in this direction.

Getting the results of this paper to work in the Cayley table model would also be interesting from the complexity-theoretic perspective; the necessary ingredients are discussed in Remark 1.2.

Lastly, we mention that extending the results of the present paper, [31], and [35] to rings beyond fields would be very interesting. In particular, working with tensors over $\mathbb{Z}/p^k\mathbb{Z}$ is an important step towards extending the results of this paper to p -groups of class 2 without restricting them to exponent p . (This is particularly important when $p = 2$, as groups of exponent 2 are abelian, so the hardest instances of 2-groups, rather than “ p -groups of class 2 and exponent p ” with $p = 2$, are often taken to be 2-groups of class 2 and exponent four.)

It seems conceivable that many of our arguments could extend to tensors over local rings—those with a unique maximal ideal—as many of our arguments are rank-based, and rank still has nice properties over local rings (e.g. Nakayama’s Lemma). In particular, if R is a ring and $\mathfrak{m}$ a maximal ideal, then $R/\mathfrak{m}$ is a field; in a local ring, there is a unique maximal ideal, so the field $R/\mathfrak{m}$ is canonically associated to R , and one can talk cleanly about rank and dimension of R -modules considered over the field $R/\mathfrak{m}$. Besides $\mathbb{Z}/p^k\mathbb{Z}$, another local ring of interest is the ring $\mathbb{F}[[t]]$ of power series in one variable over a field $\mathbb{F}$; a tensor over $\mathbb{F}[[t]]$ is essentially a 1-parameter family of tensors over $\mathbb{F}$, so studying tensor problems over $\mathbb{F}[[t]]$ could have applications to border rank and geometric complexity theory.

ACKNOWLEDGMENTS

The authors would like to thank James B. Wilson for related discussions, and Ryan Williams for pointing out the problem of distinguishing between ETH and #ETH. They would like to also thank anonymous reviewers of CCC and TOCT for careful reading, helpful feedback, and pointing out some minor gaps in some proofs. J. A. G. would like to thank V. Futorny and V. V. Sergeichuk for their collaboration on the related work [31]. Ideas leading to this work originated from the 2015 workshop “Wildness in computer science, physics, and mathematics” at the Santa Fe Institute. Both authors were supported by NSF grant DMS-1750319. Y. Q. was partly supported by Australian Research Council DP200100950.

REFERENCES

- [1] Manindra Agrawal, Neeraj Kayal, and Nitin Saxena. 2004. PRIMES is in P. *Ann. of Math. (2)* 160, 2 (2004), 781–793. <https://doi.org/10.4007/annals.2004.160.781>
- [2] Eric Allender and Bireswar Das. 2017. Zero knowledge and circuit minimization. *Inf. Comput.* 256 (2017), 2–8. <https://doi.org/10.1016/j.ic.2017.04.004>
- [3] Vikraman Arvind and Jacobo Torán. 2005. Isomorphism Testing: Perspective and Open Problems. *Bulletin of the EATCS* 86 (2005), 66–84.
- [4] László Babai. 2016. Graph isomorphism in quasipolynomial time [extended abstract]. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016*. 684–697. <https://doi.org/10.1145/2897518.2897542> arXiv:1512.03547 [cs.DS] version 2.
- [5] László Babai, Paolo Codenotti, Joshua A. Grochow, and Youming Qiao. 2011. Code equivalence and group isomorphism. In *Proceedings of the Twenty-Second Annual ACM–SIAM Symposium on Discrete Algorithms (SODA11)*. SIAM, Philadelphia, PA, 1395–1408. <https://doi.org/10.1137/1.9781611973082.107>
- [6] László Babai, Paolo Codenotti, and Youming Qiao. 2012. Polynomial-Time Isomorphism Test for Groups with No Abelian Normal Subgroups - (Extended Abstract). In *Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Proceedings, Part I*. 51–62. https://doi.org/10.1007/978-3-642-31594-7_5
- [7] László Babai and Youming Qiao. 2012. Polynomial-time Isomorphism Test for Groups with Abelian Sylow Towers. In *29th STACS*. Springer LNCS 6651, 453 – 464. <https://doi.org/10.4230/LIPIcs.STACS.2012.453>
- [8] Reinhold Baer. 1938. Groups with abelian central quotient group. *Trans. AMS* 44, 3 (1938), 357–386. <https://doi.org/10.1090/S0002-9947-1938-1501972-1>
- [9] Mihir Bellare and Shafi Goldwasser. 1994. The Complexity of Decision Versus Search. *SIAM J. Comput.* 23, 1 (1994), 97–119. <https://doi.org/10.1137/S0097539792228289>

- [10] Hans Ulrich Besche and Bettina Eick. 1999. Construction of finite groups. *J. Symb. Comput.* 27, 4 (1999), 387–404. <https://doi.org/10.1006/jsc.1998.0258>
- [11] Hans Ulrich Besche, Bettina Eick, and E.A. O’Brien. 2002. A Millennium Project: Constructing Small Groups. *Intern. J. Alg. and Comput* 12 (2002), 623–644. <https://doi.org/10.1142/S0218196702001115>
- [12] Anton Betten, Michael Braun, Harald Friepertinger, Adalbert Kerber, Axel Kohnert, and Alfred Wassermann. 2006. *Error-correcting linear codes: Classification by isometry and applications*. Vol. 18. Springer Science and Business Media.
- [13] Peter A. Brooksbank, Joshua A. Grochow, Yanan Li, Youming Qiao, and James B. Wilson. 2019. Incorporating Weisfeiler–Leman into algorithms for group isomorphism. arXiv:1905.02518 [cs.CC].
- [14] Peter A. Brooksbank, Yanan Li, Youming Qiao, and James B. Wilson. 2020. Improved Algorithms for Alternating Matrix Space Isometry: From Theory to Practice. In *28th Annual European Symposium on Algorithms, ESA 2020, September 7–9, 2020, Pisa, Italy (Virtual Conference) (LIPIcs, Vol. 173)*, Fabrizio Grandoni, Grzegorz Herman, and Peter Sanders (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 26:1–26:15. <https://doi.org/10.4230/LIPIcs.ESA.2020.26>
- [15] Peter A. Brooksbank and Eugene M. Luks. 2008. Testing isomorphism of modules. *J. Algebra* 320, 11 (2008), 4020 – 4029. <https://doi.org/10.1016/j.jalgebra.2008.07.014>
- [16] Peter A. Brooksbank, Joshua Maglione, and James B. Wilson. 2017. A fast isomorphism test for groups whose Lie algebra has genus 2. *J. Algebra* 473 (2017), 545–590. <https://doi.org/Afastisomorphismtestforgroupswhose{Lie}algebrahasgenus2>
- [17] Peter A. Brooksbank and E. A. O’Brien. 2008. Constructing the group preserving a system of forms. *Internat. J. Algebra Comput.* 18, 2 (2008), 227–241. <https://doi.org/10.1142/S021819670800441X>
- [18] John Cannon and Derek F. Holt. 2003. Automorphism group computation and isomorphism testing in finite groups. *J. Symbolic Comput.* 35, 3 (2003), 241–267.
- [19] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. 2009. Settling the complexity of computing two-player Nash equilibria. *J. ACM* 56, 3 (2009), Art. 14, 57. <https://doi.org/10.1145/1516512.1516516>
- [20] Alexander Chistov, Gábor Ivanyos, and Marek Karpinski. 1997. Polynomial time algorithms for modules over finite dimensional algebras. In *Proceedings of the 1997 International Symposium on Symbolic and Algebraic Computation (Kihei, Maui, Hawaii, USA) (ISSAC ’97)*. ACM, 68–74. <https://doi.org/10.1145/258726.258751>
- [21] Serena Cicalò, Willem A. de Graaf, and Michael Vaughan-Lee. 2012. An effective version of the Lazard correspondence. *J. Algebra* 352, 1 (2012), 430 – 450. <https://doi.org/10.1016/j.jalgebra.2011.11.031>
- [22] Henry Cohn. 2004. Projective Geometry over $\mathbb{F}_1$ and the Gaussian Binomial Coefficients. *The American Mathematical Monthly* 111, 6 (2004), 487–495.
- [23] Henry H Crapo. 1966. The Möbius function of a lattice. *Journal of Combinatorial Theory* 1, 1 (1966), 126–131.
- [24] W.A. de Graaf. 2000. *Lie Algebras: Theory and Algorithms*. North-Holland Mathematical Library, Vol. 56. Elsevier Science.
- [25] Holger Dell, Thore Husfeldt, Dániel Marx, Nina Taslamani, and Martin Wahlén. 2014. Exponential time complexity of the permanent and the Tutte polynomial. *ACM Transactions on Algorithms (TALG)* 10, 4 (2014), 1–32.
- [26] Sean Eberhard and Stefan-C. Virchow. 2020. Random generation of the special linear group. *Trans. Amer. Math. Soc.* (March 2020), 1. <https://doi.org/10.1090/tran/8009>
- [27] Wayne Eberly and Mark Giesbrecht. 2000. Efficient decomposition of associative algebras over finite fields. *Journal of Symbolic Computation* 29, 3 (2000), 441–458. <https://doi.org/10.1006/jsc.1999.0308>
- [28] Bettina Eick, C. R. Leedham-Green, and E. A. O’Brien. 2002. Constructing automorphism groups of p -groups. *Comm. Algebra* 30, 5 (2002), 2271–2295. <https://doi.org/10.1081/AGB-120003468>
- [29] V. Felsch and J. Neubüser. 1970. On a programme for the determination of the automorphism group of a finite group.. In *Computational Problems in Abstract Algebra (Proceedings of a Conference on Computational Problems in Algebra, Oxford, 1967)*, Pergamon J. Leech (Ed.). Oxford, 59–60.
- [30] Katalin Friedl and Lajos Rónyai. 1985. Polynomial Time Solutions of Some Problems in Computational Algebra. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing, May 6–8, 1985, Providence, Rhode Island, USA*, Robert Sedgewick (Ed.). ACM, 153–162. <https://doi.org/10.1145/22145.22162>
- [31] Vyacheslav Futorny, Joshua A. Grochow, and Vladimir V. Sergeichuk. 2019. Wildness for tensors. *Lin. Alg. Appl.* 566 (2019), 212–244. <https://doi.org/10.1016/j.laa.2018.12.022>
- [32] Joshua A. Grochow. 2019. Answer to “What is the hardest instance for the group isomorphism problem?” on Theoretical Computer Science StackExchange. <https://cstheory.stackexchange.com/a/42551/129>.
- [33] Joshua A. Grochow and Youming Qiao. 2015. Polynomial-Time Isomorphism Test of Groups that are Tame Extensions - (Extended Abstract). In *Algorithms and Computation - 26th International Symposium, ISAAC 2015, Nagoya, Japan, December 9–11, 2015, Proceedings*. 578–589. https://doi.org/10.1007/978-3-662-48971-0_49
- [34] Joshua A. Grochow and Youming Qiao. 2017. Algorithms for group isomorphism via group extensions and cohomology. *SIAM J. Comput.* 46, 4 (2017), 1153–1216. <https://doi.org/10.1137/15M1009767> Preliminary version in IEEE Conference on Computational Complexity (CCC) 2014 (DOI:10.1109/CCC.2014.19). Also available as arXiv:1309.1776 [cs.DS] and ECCC Technical Report TR13-123..

- [35] Joshua A. Grochow and Youming Qiao. 2021. On the complexity of isomorphism problems for tensors, groups, and polynomials I: Tensor Isomorphism-completeness. In *ITCS*. to appear. arXiv:1907.00309.
- [36] Martin Grohe and Pascal Schweitzer. 2020. The graph isomorphism problem. *Commun. ACM* 63, 11 (2020), 128–134. <https://doi.org/10.1145/3372123>
- [37] Xiaoyu He and Youming Qiao. 2020. On the Baer–Lovász–Tutte construction of groups from graphs: isomorphism types and homomorphism notions. arXiv:2003.07200 [math.CO].
- [38] Russell Impagliazzo and Ramamohan Paturi. 2001. On the complexity of k -SAT. *J. Comput. System Sci.* 62, 2 (2001), 367–375.
- [39] Gábor Ivanyos. 2000. Fast randomized algorithms for the structure of matrix algebras over finite fields. In *Proceedings of the 2000 international symposium on Symbolic and algebraic computation*. ACM, 175–183. <https://doi.org/10.1145/345542.345620>
- [40] Gábor Ivanyos, Marek Karpinski, and Nitin Saxena. 2010. Deterministic Polynomial Time Algorithms for Matrix Completion Problems. *SIAM J. Comput.* 39, 8 (2010), 3736–3751. <https://doi.org/10.1137/090781231>
- [41] Gábor Ivanyos and Lajos Rónyai. 1999. Computations in associative and Lie algebras. In *Some tapas of computer algebra*. Springer, 91–120. https://doi.org/10.1007/978-3-662-03891-8_5
- [42] Zhengfeng Ji, Youming Qiao, Fang Song, and Aaram Yun. 2019. General Linear Group Action on Tensors: A Candidate for Post-quantum Cryptography. In *Theory of Cryptography - 17th International Conference, TCC 2019, Nuremberg, Germany, December 1-5, 2019, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 11891)*, Dennis Hofheinz and Alon Rosen (Eds.). Springer, 251–281. https://doi.org/10.1007/978-3-030-36030-6_11 Preprint arXiv:1906.04330 [cs.CR].
- [43] William M. Kantor. 1990. Some topics in asymptotic group theory. *Groups, Combinatorics and Geometry (Durham)* (1990), 403–421.
- [44] William M Kantor and Alexander Lubotzky. 1990. The probability of generating a finite classical group. *Geometriae Dedicata* 36, 1 (1990), 67–87.
- [45] Neeraj Kayal and Timur Nezhmetdinov. 2009. Factoring Groups Efficiently. In *Automata, Languages and Programming, 36th International Colloquium, ICALP 2009, Rhodes, Greece, July 5-12, 2009, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 5555)*, Susanne Albers, Alberto Marchetti-Spaccamela, Yossi Matias, Sotiris E. Nikolettseas, and Wolfgang Thomas (Eds.). Springer, 585–596. https://doi.org/10.1007/978-3-642-02927-1_49 Preprint ECCC Tech. Report TR08-074.
- [46] E. I. Khukhro. 1998. p -automorphisms of finite p -groups. London Mathematical Society Lecture Note Series, Vol. 246. Cambridge University Press, Cambridge. xviii+204 pages. <https://doi.org/10.1017/CBO9780511526008>
- [47] Johannes Köbler, Uwe Schöning, and Jacobo Torán. 1993. *The graph isomorphism problem: its structural complexity*. Birkhauser Verlag, Basel, Switzerland, Switzerland. <https://doi.org/10.1007/978-1-4612-0333-9>
- [48] Tamara G Kolda and Brett W Bader. 2009. Tensor decompositions and applications. *SIAM review* 51, 3 (2009), 455–500. <https://doi.org/10.1137/07070111X>
- [49] Michel Lazard. 1954. Sur les groupes nilpotents et les anneaux de Lie. *Ann. Sci. Ecole Norm. Sup. (3)* 71 (1954), 101–190. <https://doi.org/0.24033/asens.1021>
- [50] François Le Gall. 2009. Efficient Isomorphism Testing for a Class of Group Extensions. In *Proc. 26th STACS*. 625–636. <https://doi.org/10.4230/LIPIcs.STACS.2009.1830>
- [51] Mark L. Lewis and James B. Wilson. 2012. Isomorphism in expanding families of indistinguishable groups. *Groups Complex. Cryptol.* 4, 1 (2012), 73–110. <https://doi.org/10.1515/gcc-2012-0008>
- [52] Yinan Li and Youming Qiao. 2017. Linear Algebraic Analogues of the Graph Isomorphism Problem and the Erdős–Rényi Model. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017*, Chris Umans (Ed.). IEEE Computer Society, 463–474. <https://doi.org/10.1109/FOCS.2017.49>
- [53] Eugene M. Luks. 1992. Computing in Solvable Matrix Groups. In *FOCS 1992, 33rd Annual Symposium on Foundations of Computer Science*. IEEE Computer Society, 111–120. <https://doi.org/10.1109/SFCS.1992.267813>
- [54] Eugene M. Luks. 1993. Permutation groups and polynomial-time computation. In *Groups and computation (New Brunswick, NJ, 1991)*. DIMACS Ser. Discrete Math. Theoret. Comput. Sci., Vol. 11. Amer. Math. Soc., Providence, RI, 139–175.
- [55] Eugene M. Luks. 1999. Hypergraph Isomorphism and Structural Equivalence of Boolean Functions. In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing, May 1-4, 1999, Atlanta, Georgia, USA*. 652–658. <https://doi.org/10.1145/301250.301427>
- [56] Rudolf Mathon. 1979. A note on the graph isomorphism counting problem. *Inform. Process. Lett.* 8, 3 (1979), 131–136.
- [57] Brendan D. McKay. 1980. Practical graph isomorphism. *Congr. Numer.* (1980), 45–87.
- [58] Brendan D. McKay and Adolfo Piperno. 2014. Practical graph isomorphism, II. *Journal of Symbolic Computation* 60, 0 (2014), 94 – 112. <https://doi.org/10.1016/j.jsc.2013.09.003>
- [59] Alan H. Mekler. 1981. Stability of nilpotent groups of class 2 and prime exponent. *The Journal of Symbolic Logic* 46, 4 (1981), 781–788.
- [60] Gary L. Miller. 1978. On the $n^{\log n}$ isomorphism technique (A Preliminary Report). In *STOC*. ACM, 51–58. <https://doi.org/10.1145/800133.804331>
- [61] Takunari Miyazaki. 1996. Luks’s reduction of graph isomorphism to code equivalence. Comment to E. W. Clark. <https://groups.google.com/forum/#!msg/sci.math.research/puZxGj9HXKI/CeyH2yzyNFUJ>

- [62] Vipul Naik. 2013. *Lazard correspondence up to isoclinism*. Ph.D. Dissertation. The University of Chicago. <https://vipulnaik.com/thesis/>
- [63] Jacques Patarin. 1996. Hidden Fields Equations (HFE) and Isomorphisms of Polynomials (IP): Two New Families of Asymmetric Algorithms. In *Advances in Cryptology - EUROCRYPT '96, International Conference on the Theory and Application of Cryptographic Techniques, Saragossa, Spain, May 12-16, 1996, Proceeding*. 33–48. https://doi.org/10.1007/3-540-68339-9_4
- [64] Erez Petrank and Ron M. Roth. 1997. Is code equivalence easy to decide? *IEEE Trans. Inf. Theory* 43, 5 (1997), 1602–1604. <https://doi.org/10.1109/18.623157>
- [65] Youming Qiao, Jayalal M. N. Sarma, and Bangsheng Tang. 2011. On Isomorphism Testing of Groups with Normal Hall Subgroups. In *Proc. 28th STACS*. 567–578. <https://doi.org/10.4230/LIPIcs.STACS.2011.567>
- [66] Lajos Rónyai. 1990. Computing the Structure of Finite Algebras. *J. Symb. Comput.* 9, 3 (1990), 355–373. [https://doi.org/10.1016/S0747-7171\(08\)80017-X](https://doi.org/10.1016/S0747-7171(08)80017-X)
- [67] David J. Rosenbaum. 2013. Bidirectional collision detection and faster deterministic isomorphism testing. arXiv preprint arXiv:1304.3935 [cs.DS].
- [68] David J. Rosenbaum. 2013. Breaking the $n^{\log n}$ Barrier for Solvable-Group Isomorphism. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 1054–1073. Preprint arXiv:1205.0642 [cs.DS].
- [69] Nicolas Sendrier and Dimitris E. Simos. 2013. The Hardness of Code Equivalence over $\mathbb{F}_q$ and Its Application to Code-Based Cryptography. In *International Workshop on Post-Quantum Cryptography*. Springer, 203–216.
- [70] Seinosuke Toda. 1991. PP is as Hard as the Polynomial-Time Hierarchy. *SIAM J. Comput.* 20, 5 (1991), 865–877. <https://doi.org/10.1137/0220053>
- [71] Leslie G. Valiant. 1976. Relative complexity of checking and evaluating. *Information processing letters* 5, 1 (1976), 20–23.
- [72] James Wilson. 2014. 2014 conference on *Groups, Computation, and Geometry* at Colorado State University, co-organized by P. Brooksbank, A. Hulpke, T. Penttilä, J. Wilson, and W. Kantor. Personal communication.
- [73] James B. Wilson. 2010. Finding direct product decompositions in polynomial time. arXiv:1005.0548 [math.GR].
- [74] James B. Wilson. 2012. Existence, algorithms, and asymptotics of direct product decompositions, I. *Groups Complex. Cryptol.* 4, 1 (2012), 33–72. <https://doi.org/10.1515/gcc-2012-0007>
- [75] V. N. Zemlyachenko, N. M. Korneenko, and R. I. Tyshkevich. 1985. Graph isomorphism problem. *J. Soviet Math.* 29, 4 (01 May 1985), 1426–1481. <https://doi.org/10.1007/BF02104746>