Faster Isomorphism for p-Groups of Class 2 and Exponent p

Xiaorui Sun
xiaorui@uic.edu
University of Illinois at Chicago
United States

ABSTRACT

The group isomorphism problem determines whether two groups, given by their Cayley tables, are isomorphic. For groups with order n, an algorithm with $n^{(\log n+O(1))}$ running time, attributed to Tarjan, was proposed in the 1970s (Miller, STOC 1978). Despite the extensive study over the past decades, the current best group isomorphism algorithm has an $n^{(1/4+o(1))\log n}$ running time (Rosenbaum 2013).

The isomorphism testing for p-groups of (nilpotent) class 2 and exponent p has been identified as a major barrier to obtaining an $n^{o(\log n)}$ time algorithm for the group isomorphism problem. Although the p-groups of class 2 and exponent p have much simpler algebraic structures than general groups, the best-known isomorphism testing algorithm for this group class also has an $n^{O(\log n)}$ running time.

In this paper, we present an isomorphism testing algorithm for p-groups of class 2 and exponent p with running time $n^{O((\log n)^{5/6})}$ for any prime $p > 2$. Our result is based on a novel reduction to the skew-symmetric matrix tuple isometry problem (Ivanyos and Qiao, SIAM J. Computing, 2019). To obtain the reduction, we develop several tools for matrix space analysis, including a matrix space individualization-refinement method and a characterization of the low rank matrix spaces.

CCS CONCEPTS

• **Theory of computation → Design and analysis of algorithms**;
• **Mathematics of computing → Discrete mathematics**.

KEYWORDS

group Isomorphism, p-groups of class 2 and exponent p, matrix space isometry

ACM Reference Format:
Xiaorui Sun. 2023. Faster Isomorphism for p-Groups of Class 2 and Exponent p. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC '23), June 20–23, 2023, Orlando, FL, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3564246.3585250

STOC '23, June 20–23, 2023, Orlando, FL, USA

ACM ISBN 978-1-4503-9913-5/23/06…$15.00
https://doi.org/10.1145/3564246.3585250

1 INTRODUCTION

The group isomorphism problem is to determine whether two groups, given by their Cayley (multiplication) tables, are isomorphic. The problem is among a few classes of problems in NP that are not known to be solvable in polynomial time or NP-Complete [24]. The group isomorphism problem and its variants have close connections to cryptography, computational group theory, and algebraic complexity theory [30]. Furthermore, following Babai's breakthrough on the quasi-polynomial time algorithm for graph isomorphism [4, 5], group isomorphism has become a bottleneck for an $n^{o(\log n)}$ time algorithm of graph isomorphism because group isomorphism reduces to graph isomorphism.

The group isomorphism problem has been extensively studied since the 1970s [7, 8, 11, 16–18, 21, 22, 25, 26, 32, 34–36, 38, 39, 41–45, 50, 54, 55]. A simple algorithm for group isomorphism, attributed to Tarjan, picks a generating set in one of the groups and checks for all possible images of the generating set in the other group, whether the partial correspondence extends to an isomorphism [39]. Since every group of order n has a generating set of size at most $\log_2 n$, this algorithm results in an $n^{\log_2 n+O(1)}$ running time. The current best-known algorithm for the group isomorphism problem has an $n^{(1/4+o(1))\log_2 n}$ running time [43].

It is long believed that the isomorphism testing of p-groups of class 2 and exponent p is a major bottleneck for the group isomorphism problem [7, 15, 17, 18, 35, 36, 43]. A group G is a p-group of (nilpotent) class 2 and exponent p for some prime number p if every element except the identity has an order of p, and G is not abelian but $[G,[G,G]]$ only contains the identity element, where $[G,H]$ denotes the group generated by $xyx^{-1}y^{-1}$ for all $x \in G, y \in H$.

The best-known algorithm for the isomorphism testing of p-groups of class 2 and exponent p does not have a major advantage in the running time, being $n^{O(\log_2 n)}$ [43], over the general groups, even though the structure of p-groups of class 2 and exponent p was well understood [13, 51, 54, 55], and the isomorphism testing of this group class has been studied in depth [15, 17, 18, 35, 36, 43, 46]. Hence, to develop a better algorithm for isomorphism testing of general groups, it is necessary to provide a faster algorithm for p-groups of class 2 and exponent p.

1.1 Our Result

In this paper, we present an isomorphism testing algorithm for p-groups of class 2 and exponent p with $n^{o(\log n)}$ running time for any odd prime p.

Theorem 1.1. *Let G and H be two groups of order n. If both G and H are p-groups of class 2 and exponent p for some prime number $p > 2$, then given the Cayley tables of G and H, there is an algorithm with running time $n^{O((\log n)^{5/6})}$ to determine whether G and H are isomorphic.*

Theorem 1.1 utilizes the Baer's correspondence [13], which reduces the group isomorphism problem for p-groups of class 2 and exponent p to the isometry testing problem of skew-symmetric matrix spaces.

A square matrix A is a skew-symmetric matrix if $A^T = -A$. In the isometry testing problem for skew-symmetric matrix spaces, the input consists of the linear bases of two skew-symmetric matrix spaces $\mathfrak{A}$ and $\mathfrak{B}$. The problem is to decide whether there is an isometry S from $\mathfrak{A}$ to $\mathfrak{B}$, i.e., an invertible matrix S such that $S\mathfrak{A}S^T = \mathfrak{B}$, where $S\mathfrak{A}S^T$ is the linear span of the matrices SAS^T for all the matrices $A \in \mathfrak{A}$. We prove the following result for the isometry testing problem of skew-symmetric matrix spaces.

Theorem 1.2. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be two linear matrix spaces, both of dimension m, such that every matrix in $\mathfrak{A}$ or $\mathfrak{B}$ is an $n \times n$ skew-symmetric matrix over $\mathbb{F}_p$ for some prime number $p > 2$ and positive integers m, n. There is an algorithm with running time $p^{O((n+m)^{1.8} \cdot \log(p))}$ to determine whether there is an invertible $n \times n$ matrix S over $\mathbb{F}_p$ such that $S\mathfrak{A}S^T = \mathfrak{B}$.*

We obtain Theorem 1.2 by combining several new tools to analyze matrix spaces, including an individualization-refinement method for matrix spaces, a characterization of low rank matrix spaces, and a reduction from the isometry testing of skew-symmetric matrix spaces to the isometry testing of skew-symmetric matrix tuples [30].

To obtain Theorem 1.1, let k denote $\log_p(n)$. We apply Theorem 1.2 for the case of $k > (\log_2(p))^5$ by constructing the skew-symmetric matrix spaces for both input groups according to the Baer's correspondence [13]. Theorem 1.2 implies the running time for this case is $n^{O((\log n)^{5/6})}$. For the case of $k \leq (\log_2(p))^5$, we run the aforementioned generating set enumeration algorithm [39]. Because every p group of order p^k has a generating set of size at most k, the running time of the algorithm for this case is $p^{O(k^2)}$, which is also $n^{O((\log n)^{5/6})}$.

1.2 Related Work

The group isomorphism problem has been studied for variant group classes. Polynomial time algorithms have been developed for abelian groups [32, 45, 50], groups formed by semidirect products of an abelian group and a cyclic group [34, 54, 55], groups with normal Hall subgroups [42], groups with abelian Sylow towers [11], and groups with no abelian normal subgroups [8]. Dietrich and Wilson recently showed that the group isomorphism problem can be solved in nearly linear time for most orders [21].

For p-groups of class 2 and exponent p, algorithms for some nontrivial subclasses of this group class have been proposed [17, 18, 35]. Li and Qiao showed that if the p-groups of class 2 and exponent p are generated randomly, then the isomorphism testing problem can be solved in polynomial time in the average case [36]. In [15], the average case running time was further improved to linear. In this work, we focus on the isomorphism testing for p-groups of class 2 and exponent p in the worst case.

The refinement methods, such as the naive refinement [9] and Weisfeiler-Leman refinement [52], have been powerful tools for the graph isomorphism problem. The refinement methods have been successfully used for graph isomorphism testing algorithms [2, 3, 6, 9, 10, 12, 19, 20, 27–29, 33, 37, 40, 47, 48, 53, 56], including the celebrated quasi-polynomial time algorithm for graph isomorphism [4, 5].

The refinement approach does not extend to groups in a naive way. Several representations of groups that allow refinement have been proposed. In [15], the authors defined a hypergraph using recursively refinable filters and proposed applying the Weisfeiler-Leman refinement on the hypergraph. Brachter and Schweitzer proposed defining colors of group element tuples by group operation patterns of the elements involved in the tuple and applying the Weisfeiler-Leman refinement to refine the colors of element tuples [14]. Both approaches can distinguish between several non-isomorphic constructions of p-groups of class 2 and exponent p. However, it was unclear how these refinement methods could be used to develop faster worst case isomorphism testing algorithms.

The isometry testing of skew-symmetric matrix spaces was studied in [16, 25, 26, 36]. Its applications in cryptography were investigated in [15, 31, 49].

2 PRELIMINARIES

We give the notations and previous results used in this paper.

2.1 Notations

Throughout the paper, the vectors and matrices are over $\mathbb{F}_p$ for a prime number $p > 2$. Let $\mathbb{F}_p^n$ be the linear space of row vectors of length n over $\mathbb{F}_p$. Unless specified, the vectors are row vectors. We use $\langle \cdot \rangle$ to denote the linear span.

Matrices. Let $M(n, \mathbb{F}_p)$ (and respectively $M(m, n, \mathbb{F}_p)$) be the linear space of $n \times n$ (and respectively $m \times n$) matrices over $\mathbb{F}_p$. Let $\mathrm{GL}(n, \mathbb{F}_p)$ be the group of $n \times n$ invertible matrices over $\mathbb{F}_p$.

For a matrix $A \in M(m, n, \mathbb{F}_p)$, let $\mathrm{rank}(A)$ be the rank of A, and A^T be the transpose of A. A square matrix $A \in M(n, \mathbb{F}_p)$ is a skew-symmetric matrix if and only if $A = -A^T$. For any $1 \leq i \leq m, 1 \leq j \leq n$, let $A[i, j]$ be the entry of A in the i-th row and j-th column.

For two matrices $A, B \in M(m \times n, \mathbb{F}_p)$, A is lexically smaller than B if there exist $1 \leq q \leq m$ and $1 \leq r \leq n$ such that the following conditions hold:

- $A[i, j] = B[i, j]$ for any $1 \leq i \leq q-1, 1 \leq j \leq n$ or any $i = q, 1 \leq j < r$;
- $A[q, r] < B[q, r]$.

Matrix Tuples and Matrix Spaces. An $m \times n$ matrix tuple $\mathcal{A}$ of length k, denoted as $\mathcal{A} = (A_1, \ldots, A_k)$, is an element in $M(m, n, \mathbb{F}_p)^k$. For any $P \in M(\alpha, m, \mathbb{F}_p)$ and $Q \in M(n, \beta, \mathbb{F}_p)$ with some positive integers α and β, let

$$P\mathcal{A}Q := (PA_1Q, PA_2, Q, \ldots, PA_kQ).$$

An $m \times n$ matrix space $\mathfrak{A}$ is a linear subspace of $M(m, n, \mathbb{F}_p)$. For any $P \in M(\alpha, m, \mathbb{F}_p)$ and $Q \in M(n, \beta, \mathbb{F}_p)$ with some positive integers α and β, let

$$P\mathfrak{A}Q := \langle PAQ : \forall A \in \mathfrak{A} \rangle.$$

Since any linear combination of skew-symmetric matrices of the same dimension is also a skew-symmetric matrix, we use $\mathrm{SS}(n, \mathbb{F}_p)$ to denote the linear space of all the $n \times n$ skew-symmetric matrices.

2.2 Isometry and Equivalence for Matrix Tuples and Spaces

We define equivalence relations for matrix tuples.

Definition 2.1 (Matrix tuple equivalence). Let $\mathcal{A} = (A_1, \ldots, A_k)$ and $\mathcal{B} = (B_1, \ldots, B_k)$ be two matrix tuples in $M(m, n, \mathbb{F}_p)^k$. $\mathcal{A}$ and $\mathcal{B}$ are equivalent if there exist two matrices $P \in \mathrm{GL}(m, \mathbb{F}_p)$ and $Q \in \mathrm{GL}(n, \mathbb{F}_p)$ such that $P\mathcal{A}Q = \mathcal{B}$.

Definition 2.2 (Skew-symmetric matrix tuple isometry). Let $\mathcal{A} = (A_1, \ldots, A_k)$ and $\mathcal{B} = (B_1, \ldots, B_k)$ be two skew-symmetric matrix tuples in $\mathrm{SS}(n, \mathbb{F}_p)^k$. $\mathcal{A}$ and $\mathcal{B}$ are isometric if there exists a matrix $P \in \mathrm{GL}(n, \mathbb{F}_p)$ such that $P\mathcal{A}P^T = \mathcal{B}$. Such a matrix P is called an isometry from $\mathcal{A}$ to $\mathcal{B}$.

In this paper, we use the algorithm for the isometry testing of two skew-symmetric matrix tuples (Theorem 2.4) and the algorithm for the equivalence testing of two matrix tuples (Theorem 2.3), both proposed by Ivanyos and Qiao in [30].

Theorem 2.3 (Proposition 3.2 of [30]). *Given two matrix tuples $\mathcal{A} = (A_1, \ldots, A_k)$ and $\mathcal{B} = (B_1, \ldots, B_k)$ in $M(m, n, \mathbb{F}_p)^k$ for some prime $p > 2$ and positive integers k, m and n, there is an algorithm with running time $\mathrm{poly}(k, n, m, p)$ to determine whether $\mathcal{A}$ and $\mathcal{B}$ are equivalent.*

Theorem 2.4 (Theorem 1.7 of [30]). *Let $\mathcal{A} = (A_1, \ldots, A_k)$ and $\mathcal{B} = (B_1, \ldots, B_k)$ be two skew-symmetric matrix tuples of length k such that the matrices in $\mathcal{A}$ and $\mathcal{B}$ are of dimension $n \times n$ over $\mathbb{F}_p$ for some prime $p > 2$. There is an algorithm with running time $\mathrm{poly}(n, k, p)$ to determine whether there is an isometry from $\mathcal{A}$ to $\mathcal{B}$. If yes, the algorithm also returns an isometry from $\mathcal{A}$ to $\mathcal{B}$.*

Following the definitions for matrix tuples, we also define the isometry of skew-symmetric matrix spaces.

Definition 2.5 (Skew-symmetric matrix space isometry). Let $\mathfrak{A}, \mathfrak{B}$ be two skew-symmetric matrix spaces. $\mathfrak{A}$ and $\mathfrak{B}$ are isometric if there exists a matrix $P \in \mathrm{GL}(n, \mathbb{F}_p)$ such that $P\mathfrak{A}P^T = \mathfrak{B}$. P is called an isometry from $\mathfrak{A}$ to $\mathfrak{B}$ if P exists.

In Section 2.3, we will use the Baer's correspondence to reduce the group isomorphism for p-groups of class two and exponent p to the problem of isometry testing for skew-symmetric matrix spaces.

2.3 Baer's Correspondence

For a p-group of nilpotent class 2 and exponent p, let p^k denote the order of the group. Because of the class two and exponent p condition, $G/Z(G)$ is isomorphic to $\mathbb{Z}_p^n$, and $[G, G]$ is isomorphic to $\mathbb{Z}_p^m$ for some positive integers n and m such that $n + m \leq k$, where $Z(G)$ denotes the center of G, and $[G, G]$ denotes the group generated by $xyx^{-1}y^{-1}$ for all $x, y \in G$. Taking an arbitrary basis of $G/Z(G)$, an arbitrary basis of $[G, G]$, and taking the commutator bracket, we obtain a skew-symmetric bilinear map $b_G : \mathbb{F}_p^n \times \mathbb{F}_p^n \to \mathbb{F}_p^m$, which can be represented by a skew-symmetric matrix tuple $\mathcal{G} = (G_1, \ldots, G_m)$ such that every G_i is a matrix in $\mathrm{SS}(n, \mathbb{F}_p)$. Such a skew-symmetric matrix tuple is called a skew-symmetric matrix tuple of G.

For two p-groups G and H of nilpotent class 2 and exponent p, it is necessary for H to be isomorphic to G that

$$\dim_{\mathbb{Z}_p}(G/Z(G)) = \dim_{\mathbb{Z}_p}(H/Z(H))$$

and

$$\dim_{\mathbb{Z}_p}([G, G]) = \dim_{\mathbb{Z}_p}([H, H])$$

The following theorem, also called Baer's correspondence, was proved by Baer in [13].

Theorem 2.6 (Baer's correspondence [13], rephrased). *Let G and H be two p-groups of class two and exponent p for some prime number p with the same order. Let $\mathcal{G}$ and $\mathcal{H}$ be the skew-symmetric matrix tuples of G and H, respectively. If both $\mathcal{G}$ and $\mathcal{H}$ are $n \times n$ skew-symmetric matrix tuples of length m, then G and H are isomorphic if and only if there are matrices $P \in \mathrm{GL}(n, \mathbb{F}_p)$ and $Q \in \mathrm{GL}(m, \mathbb{F}_p)$ such that*

$$G_i = \sum_{j=1}^{m} Q[i, j](P \cdot H_j \cdot P^T).$$

Furthermore, we can also represent skew-symmetric matrix tuples of groups by skew-symmetric matrix spaces. Given an arbitrary skew-symmetric matrix tuple $\mathcal{G}$ of group G, the skew-symmetric matrix space $\mathfrak{G}$ of G is the linear matrix space spanned by matrices in $\mathcal{G}$. Hence, Baer's correspondence can be rephrased as follows.

Corollary 2.7. *Let G and H be two p-groups of class two and exponent p for some prime number p with the same order. Let $\mathfrak{G}$ and $\mathfrak{H}$ be the skew-symmetric matrix spaces of G and H, respectively. G and H are isomorphic if and only if $\mathfrak{G}$ and $\mathfrak{H}$ are isometric.*

3 TECHNICAL OVERVIEW

We provide an overview of the algorithm for the isometry testing of skew-symmetric matrix spaces (Theorem 1.2).

The main idea of proving Theorem 1.2 is to reduce the isometry testing of skew-symmetric matrix spaces to the isometry testing of skew-symmetric matrix tuples, which can be solved efficiently by Theorem 2.4. The difference between the two problems is that the correspondence between matrices from two matrix tuples is fixed by the indices of the matrices, whereas for matrix spaces, no such correspondence is given.

For an input of the skew-symmetric matrix space isometry testing problem, let m be the dimension of the input matrix spaces and n be the number of rows (or columns) for each square matrix in the matrix spaces. If one enumerates all the possible correspondences between the matrices of the two matrix spaces, then the running time of the algorithms is $O(p^{m^2} \cdot \mathrm{poly}(m, n, p))$, which offers no improvement on the isomorphism of p-groups of class 2 and exponent p.

In this work, we give a $p^{O((n+m)^{1.8} \cdot \log p)}$ time reduction to the skew-symmetric matrix tuple isometry problem. The reduction is obtained by investigating several new tools for analyzing the structure of skew-symmetric matrix spaces, including a matrix space individualization-refinement method (Section 3.1) and a characterization of the low rank matrix spaces (Section 3.2).

3.1 Individualization-Refinement for Matrix Spaces

One powerful technique for graph isomorphism is the method of individualization-refinement [2–4, 6, 9, 12, 19, 47, 48, 56]. For graphs, the individualization-refinement method first chooses a set of a small number of vertices and assigns each chosen vertex

a distinct vertex color, and then it refines the vertex colors by assigning distinguished vertices different colors in a canonical way until vertices of the same color cannot be further distinguished.

A natural question for the group isomorphism problem is whether it is possible to define individualization-refinement operations for group isomorphism. Based on the connection between group isomorphism for p-groups of class 2 and exponent p and the skew-symmetric matrix space isometry problem [13], Li and Qiao proposed a matrix space individualization-refinement method, which follows the individualization-refinement for random graphs [9], and analyzed the isometry testing of skew-symmetric matrix spaces in the average case [36].

In this work, we propose a different matrix space individualization-refinement to enable the analysis of the isometry of skew-symmetric matrix spaces in the worst case. Consider an $m \times n$ matrix space $\mathfrak{A}$. The individualization in our scenario is defined by a left individualization matrix L and a right individualization matrix R, where L is a matrix with m columns and R is a matrix with n rows. In the refinement, we compute LAR for each matrix $A \in \mathfrak{A}$. If $LA'R$ does not equal $LA''R$ for some $A', A'' \in \mathfrak{A}$, then A' and A'' are distinguished.

Ideally, if $LA'R$ does not equal $LA''R$ for any two matrices A' and A'' in $\mathfrak{A}$, then each matrix A in the space can be uniquely identified by LAR, and thus all the matrices in $\mathfrak{A}$ are distinguished. Consider two isometric skew-symmetric matrix spaces $\mathfrak{A}$ and $\mathfrak{B}$. Let $L_{\mathfrak{A}}$ and $R_{\mathfrak{A}}$ be individualization matrices for $\mathfrak{A}$ that distinguish all the matrices in $\mathfrak{A}$. Let $L_{\mathfrak{B}}$ and $R_{\mathfrak{B}}$ be individualization matrices for $\mathfrak{B}$ such that

$$L_{\mathfrak{B}} = L_{\mathfrak{A}} S^{-1} \text{ and } R_{\mathfrak{B}} = \left(S^T\right)^{-1} R_{\mathfrak{A}}$$

for some isometry S from $\mathfrak{A}$ to $\mathfrak{B}$. One can distinguish all the matrices in both spaces by their individualization matrices and then establish a bijection between the matrices in the two spaces. Thus the skew-symmetric matrix space isometry problem reduces to the skew-symmetric matrix tuple isometry problem, which can be efficiently solved by Theorem 2.4. Furthermore, suppose $L_{\mathfrak{A}}$ contains a small number of rows and $R_{\mathfrak{A}}$ contains a small number of columns. Then one can solve the skew-symmetric matrix space isometry problem efficiently by enumerating all the possible corresponding $L_{\mathfrak{B}}$ and $R_{\mathfrak{B}}$.

We show that the number of rows for the left individualization matrices and the number of columns for the right individualization matrices are related to the rank of matrices in the matrix space.

Lemma 3.1. *Let $\mathfrak{A}$ be a d-dimensional matrix subspace of $M(m, n, \mathbb{F}_p)$ for a prime p and some positive integers d, m, n. For any $k \geq 4$, denote*

$$t := \left\lceil 32 \max\{d \log(p), k\} / \sqrt{k} \right\rceil.$$

There is a left individualization matrix $L \in M(t, m, \mathbb{F}_p)$ and a right individualization matrix $R \in M(n, t, \mathbb{F}_p)$ such that for any $A \in \mathfrak{A}$ of rank at least k, LAR is a non-zero matrix.

By Lemma 3.1, if every matrix (except the zero matrix) in a skew-symmetric matrix space is of high rank, then the skew-symmetric matrix space isometry problem reduces to the skew-symmetric matrix tuple isometry problem efficiently.

3.2 Low Rank Matrix Space Characterization

The hard case for the matrix space individualization/refine method is that there are some matrices A in the space such that LAR are zero matrices. Because of the linearity, such matrices form a linear subspace of the original matrix space. To tackle this hard case, we characterize the structure of the matrix space in which every matrix is of low rank. Such a matrix space is called a low rank matrix space.

As our main technical result for the low rank matrix space characterization, we show that, for a matrix space $\mathfrak{A}$ such that every matrix in the space is of rank at most r, there are invertible matrices P and Q, called left and right formatting matrices, such that for each $A \in \mathfrak{A}$, PAQ has non-zero entries only in the last $O(r^2)$ rows or columns.

Lemma 3.2. *Let $\mathfrak{A}$ be a matrix subspace of $M(m, n, \mathbb{F}_p)$ or a skew-symmetric matrix subspace of $\mathrm{SS}(n, \mathbb{F}_p)$ such that for each $A \in \mathfrak{A}$, $\mathrm{rank}(A) \leq r$ for some positive integer r. There is a matrix $P \in \mathrm{GL}(m, \mathbb{F}_p)$, a matrix $Q \in \mathrm{GL}(n, \mathbb{F}_p)$, and an integer $\ell = O(r^2)$ such that for any $A \in \mathfrak{A}$, $PAQ[i, j] = 0$ for all the $1 \leq i \leq n - \ell$ and $1 \leq j \leq n - \ell$. Furthermore, $Q = P^T$ if $\mathfrak{A}$ is a skew-symmetric matrix space.*

We remark that similar characterizations were studied in [1, 23]. But to the author's knowledge, all the previous results require that the underlying field has at least $r + 1$ elements.

Together with matrix space individualization-refinement, we can represent a matrix space in a more structured way. First, we construct a "semi-canonical" basis for the input matrix space. Suppose we apply left and right individualization matrices L and R to a matrix space $\mathfrak{A}$ of dimension d and compute a linear basis $(A_1, \ldots, A_d)$ of $\mathfrak{A}$ such that $(LA_1R, LA_2R, \ldots, LA_dR)$ is lexically minimized among all the linear basis of $\mathfrak{A}$. Because the zero matrix is lexically the smallest among all the matrices, the first few matrices in the semi-canonical basis correspond to a linear basis of $\mathfrak{C}$, which is the linear span of all the matrices $A \in \mathfrak{A}$ such that LAR is a zero matrix.

We further apply formatting matrices P and Q for $\mathfrak{C}$ to each matrix in the semi-canonical basis of $\mathfrak{A}$ (every matrix A in the semi-canonical basis becomes PAQ). Then by our low rank matrix space characterization, the matrices that form a linear basis of $\mathfrak{C}$ have non-zero entries only in the last few rows or columns. See Figure 1 for an illustration.

The semi-canonical basis is not canonical because, for fixed individualization matrices, there can be different semi-canonical bases. But the semi-canonical bases can provide a partial correspondence between two isometric skew-symmetric matrix spaces. Suppose two skew-symmetric matrix spaces $\mathfrak{A}$ and $\mathfrak{B}$ are isometric and let S be an isometry from $\mathfrak{A}$ to $\mathfrak{B}$. For individualization matrices L and R of $\mathfrak{A}$, let $(A_1, \ldots, A_d)$ be a semi-canonical basis of $\mathfrak{A}$ with L and R as individualization matrices, and $(B_1, \ldots, B_d)$ be a semi-canonical basis of $\mathfrak{B}$ with LS^{-1} and $(S^T)^{-1}R$ as individualization matrices. Then for each $1 \leq i \leq d$,

$$SA_iS^T = B_i + B_i'$$

for some B_i' satisfying the condition that

$$LS^{-1}B_i' \left(S^T\right)^{-1} R$$

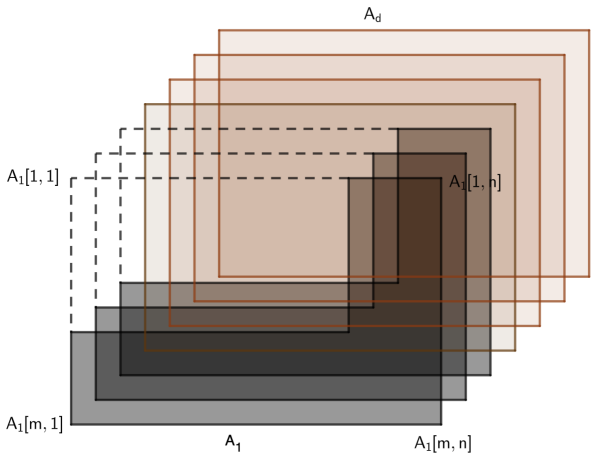

Figure 1: The semi-canonical basis of a matrix space after applying matrix space individualization-refinement and the low rank matrix space characterization. The three black matrices in the front form a basis of the space spanned by all the matrices $A \in \mathfrak{A}$ such that LAR is a zero matrix. The transparent rectangles enclosed by the dashed black lines are zero matrices. The four light brown matrices in the back are the rest matrices in the basis.

is a zero matrix. The partial correspondence also holds for two equivalent matrix spaces. Two matrix spaces $\mathfrak{A}$ and $\mathfrak{B}$, in which matrices are not necessarily square matrices, are equivalent if there are invertible matrices X and Y such that $X\mathfrak{A}Y = \mathfrak{B}$, i.e., $\mathfrak{B}$ equals the space spanned by XAY for all the matrices $A \in \mathfrak{A}$.

3.3 Tensor Representation of Skew-Symmetric Matrix Spaces

Next, we combine the matrix space individualization-refinement and the low rank matrix space characterization to analyze skew-symmetric matrix spaces. For convenience, let us define a three-tensor representation for skew-symmetric matrix spaces following [36]. For a skew-symmetric matrix space $\mathfrak{A}$ of dimension m such that every matrix in the space is an $n \times n$ matrix, a three-tensor $\mathbf{G} \in \mathbb{F}_p^{m \times n \times n}$ is a skew-symmetric matrix space tensor of $\mathfrak{A}$ if $\mathbf{G}[i, j, k] = A_i[j, k]$ for a linear basis $(A_1, \ldots, A_m)$ of $\mathfrak{A}$, where $A_i[j, k]$ is the (j, k)-th entry of A_i, and $\mathbf{G}[i, j, k]$ is the (i, j, k)-th entry of $\mathbf{G}$.

Given a skew-symmetric matrix space tensor $\mathbf{G}$, we use $\mathfrak{X}_{\mathbf{G},i}$ to denote the $n \times n$ skew-symmetric matrix such that

$$\mathfrak{X}_{\mathbf{G},i}[j, k] = \mathbf{G}[i, j, k],$$

use $\mathfrak{Y}_{\mathbf{G},j}$ to denote the $m \times n$ matrix such that

$$\mathfrak{Y}_{\mathbf{G},j}[i, k] = \mathbf{G}[i, j, k],$$

and use $\mathfrak{Z}_{\mathbf{G},k}$ to denote the $m \times n$ matrix such that

$$\mathfrak{Z}_{\mathbf{G},k}[i, j] = \mathbf{G}[i, j, k].$$

Since A_i are skew-symmetric matrices for all the $1 \le i \le m$, $\mathfrak{X}_{\mathbf{G},i}$ are skew symmetric matrices for all the $1 \le i \le m$, and $\mathfrak{Y}_{\mathbf{G},j}$ equals to $-\mathfrak{Z}_{\mathbf{G},j}$ for all the $1 \le j \le n$.

We also use $\mathfrak{X}_{\mathbf{G}}$ to denote the matrix space $\langle \mathfrak{X}_{\mathbf{G},1}, \ldots \mathfrak{X}_{\mathbf{G},m} \rangle$, use $\mathfrak{Y}_{\mathbf{G}}$ to denote the matrix space $\langle \mathfrak{Y}_{\mathbf{G},1}, \ldots \mathfrak{Y}_{\mathbf{G},n} \rangle$, and use $\mathfrak{Z}_{\mathbf{G}}$ to denote the matrix space $\langle \mathfrak{Z}_{\mathbf{G},1}, \ldots \mathfrak{Z}_{\mathbf{G},n} \rangle$, where $\langle \cdot \rangle$ is the linear span. We remark that $\mathfrak{X}_{\mathbf{G}}$ is a skew-symmetric matrix space, but $\mathfrak{Y}_{\mathbf{G}}$ and $\mathfrak{Z}_{\mathbf{G}}$ are not.

One can verify that two skew-symmetric matrix spaces are isometric if and only if their tensors (denoted as $\mathbf{G}$ and $\mathbf{H}$) are isometric, i.e., there is an $n \times n$ invertible matrix N and an $m \times m$ invertible matrix M such that the transform of $\mathbf{G}$ by N and M, denoted as $\mathrm{Trans}_{N,M}(\mathbf{G})$, equals $\mathbf{H}$, where

$$\mathfrak{X}_{\mathrm{Trans}_{N,M}(\mathbf{G}),i} = \sum_{i'=1}^{m} M[i, i'] \cdot \left(N \cdot \mathfrak{X}_{\mathbf{G},i'} \cdot N^T \right).$$

3.4 Semi-Canonical Form of Skew-Symmetric Matrix Space Tensors

The purpose of the tensor representation of a skew-symmetric matrix space is to incorporate the matrix space individualization-refinement and the low rank matrix space characterization techniques so the tensor is transformed into a more structured form, called the "semi-canonical form" of the tensor.

For a skew-symmetric matrix space tensor $\mathbf{G}$, the semi-canonical form of $\mathbf{G}$, denoted as $\mathrm{SC}(\mathbf{G})$, is obtained by applying the two techniques to the three matrix spaces $\mathfrak{X}_{\mathbf{G}}$, $\mathfrak{Y}_{\mathbf{G}}$, and $\mathfrak{Z}_{\mathbf{G}}$ so matrices in each of the three matrix spaces have the structure shown in Figure 1. To achieve this, we need to carefully choose the individualization and formatting matrices in a coordinated fashion. In particular, if the individualization and formatting matrices are chosen such that, for the left formatting matrix P used for $\mathfrak{X}_{\mathbf{G}}$, P^T can also be used as the right formatting matrix for $\mathfrak{Y}_{\mathbf{G}}$ and $\mathfrak{Z}_{\mathbf{G}}$, then the tensor semi-canonical form has the structure shown in Figure 2(a). The tensor values in the transparent region are all zero. The union of the transparent region and the red cube is called the kernel of the tensor semi-canonical form. The blue region is called the surface of the tensor semi-canonical form.

For fixed individualization and formatting matrices, the tensor semi-canonical form is also fixed. However, for an efficient tensor isometry testing algorithm, it is unacceptable to enumerate all possible formatting matrices, though it is affordable to enumerate all possible individualization matrices. To address this issue, we show that if the individualization matrices are fixed, and the formatting matrices are partially fixed (i.e., only a few key rows are fixed, and all the other rows satisfy certain conditions), then the kernel is fixed. This is also the reason for the term "semi-canonical form": the semi-canonical form is not unique for fixed individualization matrices and partially fixed formatting matrices, but the kernel is unique.

In other words, if two tensors are isometric, and one constructs the semi-canonical forms of the two tensors using individualization matrices and partially fixed formatting matrices that are the same up to some isometry, then the kernels of the two semi-canonical forms are identical. Therefore, to determine whether the two tensors are isometric, one only needs to check further if there are formatting

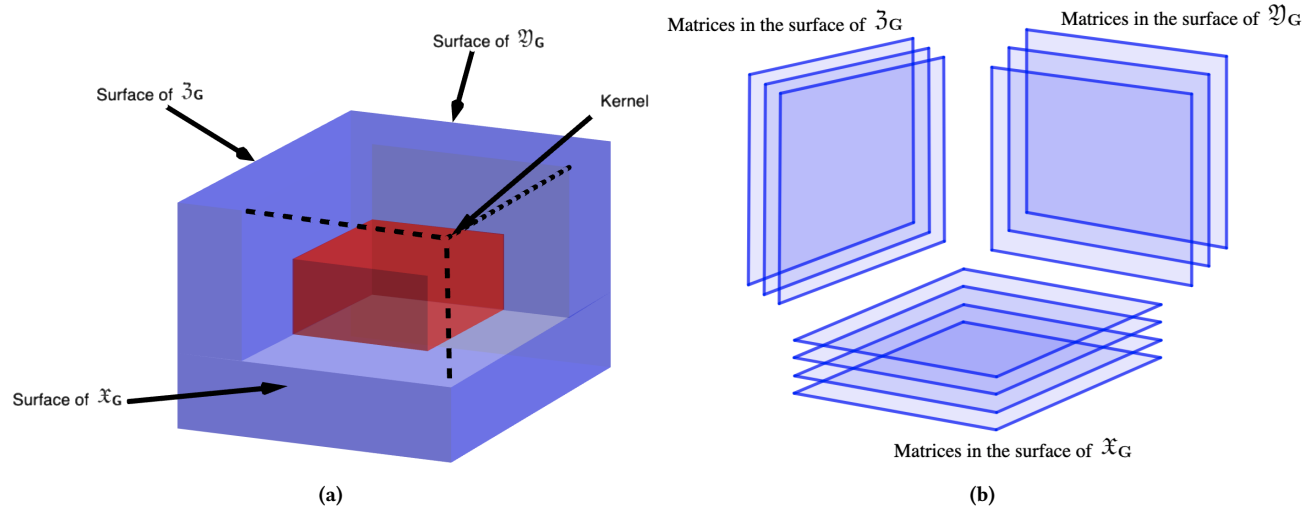

Figure 2: (a). Semi-canonical form of a skew-symmetric matrix space tensor. (b) Matrices in the surfaces of $\mathfrak{X}_G, \mathfrak{Y}_G$, and $\mathfrak{Z}_G$.

matrices that make the surface identical for the two tensors while keeping the kernel unchanged.

In addition, based on the the matrix space individualization-refinement and the low rank matrix space characterization, there are always semi-canonical forms such that the numbers of matrices in the surfaces of $\mathfrak{X}_G$, $\mathfrak{Y}_G$, and $\mathfrak{Z}_G$ (Figure 2(b)) are small. Hence, in the partially fixed formatting matrices, we also fix the rows related to the surfaces of $\mathfrak{X}_G$, $\mathfrak{Y}_G$, and $\mathfrak{Z}_G$. Then matrices in the surfaces from the three matrix spaces are fixed up to some formatting matrices satisfying the partially fixed constraint.

Hence, the isometry testing of skew-symmetric matrix space tensors reduces to the isometry testing of their semi-canonical forms by enumerating individualization matrices and partially fixed formatting matrices for both tensors. Due to the fixed kernel for all the semi-canonical forms, the isometry testing of semi-canonical forms further reduces to deciding whether the surfaces are identical between semi-canonical forms up to some formatting matrices satisfying the partially fixed constraint.

3.5 Reduction to Skew-Symmetric Matrix Tuple Isometry Testing

Finally, we reduce the isometry testing of semi-canonical forms of skew-symmetric matrix spaces to the aforementioned skew-symmetric matrix tuple isometry problem. The high-level idea is to construct a skew-symmetric matrix tuple to encode the surface of the tensor semi-canonical form. Because the matrices in the surfaces of $\mathfrak{X}_G$, $\mathfrak{Y}_G$, and $\mathfrak{Z}_G$ are fixed, we can use different matrices in the matrix tuple to encode the matrices in the surface.

Suppose the kernel is of dimension $m' \times n' \times n'$ for some $1 \le m' \le m$ and $1 \le n' \le n$. In our skew-symmetric matrix tuple of $\mathbf{SC(G)}$, denoted as $\mathcal{F}_{\mathrm{SC(G)}}$, each matrix is of dimension $(3+n+m') \times (3+n+m')$. The rows from the fourth to the $(3+n)$-th of matrices in $\mathcal{F}_{\mathrm{SC(G)}}$ correspond to the rows of matrices in $\mathfrak{X}_G$. The last m' rows of matrices in $\mathcal{F}_{\mathrm{SC(G)}}$ correspond to the first m' rows of matrices in $\mathfrak{Y}_G$ (or equivalently $\mathfrak{Z}_G$). The first three rows of matrices in $\mathcal{F}_{\mathrm{SC(G)}}$ are auxiliary rows used to ensure that the other rows satisfy the constraints of the partially fixed formatting matrices. See Figure 3 for an illustration.

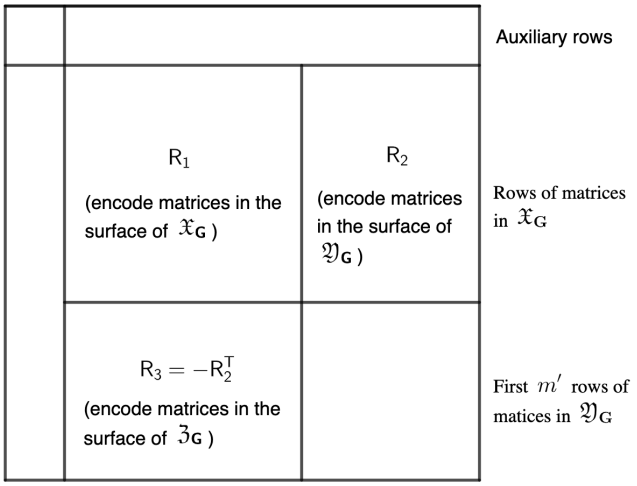

Figure 3: The matrices in $\mathcal{F}_{\mathrm{SC(G)}}$.

We use the submatrices on R_1 (as Figure 3) for all the matrices in $\mathcal{F}_{\mathrm{SC(G)}}$ to encode the skew-symmetric matrices in the surface of $\mathfrak{X}_G$. We also use the submatrices on R_2 for all the matrices in $\mathcal{F}_{\mathrm{SC(G)}}$ to encode the matrices in the surface of $\mathfrak{Y}_G$ (excluding the intersection with the surface of $\mathfrak{X}_G$). Consequently, the submatrices on R_3 for all the matrices in $\mathcal{F}_{\mathrm{SC(G)}}$, which is the negative transpose of submatrices on R_2 by the skew-symmetric condition, encode the matrices in the surface of $\mathfrak{Z}_G$ (excluding the intersection with the surface of $\mathfrak{X}_G$). We use the other submatrices to ensure constraints given by the partially fixed formatting matrices.

By carefully designing matrix tuples constructed from tensor semi-canonical forms, we show that the semi-canonical forms of two skew-symmetric matrix space tensors are isometric if and only if there is an isometry S between the skew-symmetric matrix tuples such that S is a block diagonal matrix

$$S = \begin{pmatrix} Q & 0 \\ 0 & W \end{pmatrix}$$

for some $(3+n) \times (3+n)$ matrix Q and $m' \times m'$ matrix W.

Naturally, we want to determine the isometry of the two tensors by running the skew-symmetric matrix tuple isometry algorithm (Theorem 2.4) on the matrix tuples constructed from the semi-canonical forms. However, the requirement of S being block diagonal makes things more complex.

Suppose we run the algorithm for skew-symmetric matrix tuple isometry on the matrix tuples constructed. If the algorithm returns no, then the two semi-canonical forms are not isometric. If the algorithm returns yes and an isometry that is block diagonal, then the two semi-canonical forms are isometric. The difficult case is when the algorithm returns yes and an isometry that is not block diagonal. For this case, we can neither certify that the two semi-canonical forms are isometric, nor show that the two semi-canonical forms are not isometric.

Let us consider an easier scenario: Suppose for each non-zero row vector $v \in \mathbb{F}_p^n$, there is a matrix X in the surface of $\mathfrak{X}_G$ such that vX is a non-zero vector. With this condition, together with our construction of matrix tuples, we can show that the isometry returned is of the form

$$\begin{pmatrix} X & Y \\ 0 & Z \end{pmatrix}.$$

After carefully analyzing the matrix tuples constructed, we show that

$$\begin{pmatrix} X & 0 \\ 0 & Z \end{pmatrix}$$

is also an isometry, and thus the two semi-canonical forms are isometric.

The general case is more complex because the left bottom submatrix of the isometry returned can be non-zero. However, we show that either we can certify that there exists another block diagonal isometry for the skew-symmetric matrix tuples, or we can reduce the problem to a matrix tuple equivalence problem (Definition 2.1). By Theorem 2.3, the matrix tuple equivalence problem can be solved efficiently.

ACKNOWLEDGEMENTS

Xiaorui Sun is supported by the National Science Foundation (NSF) under Grant No. 2240024 and a start-up fund from the University of Illinois at Chicago.

REFERENCES

[1] MD Atkinson and S Lloyd. 1981. Primitive spaces of matrices of bounded rank. *Journal of the Australian Mathematical Society* 30, 4 (1981), 473–482. https://doi.org/10.1017/S144678870001795X

[2] László Babai. 1980. On the complexity of canonical labeling of strongly regular graphs. *SIAM J. Comput.* 9, 1 (1980), 212–216. https://doi.org/10.1137/0209018

[3] László Babai. 1981. On the order of uniprimitive permutation groups. *Annals of Mathematics* 113, 3 (1981), 553–568. https://doi.org/10.2307/2006997

[4] László Babai. 2016. Graph isomorphism in quasipolynomial time. In *ACM Symposium on Theory of Computing (STOC)*. 684–697. https://doi.org/10.1145/2897518.2897542

[5] László Babai. 2019. Canonical form for graphs in quasipolynomial time: preliminary report. In *ACM Symposium on Theory of Computing (STOC)*. 1237–1246. https://doi.org/10.1145/3313276.3316356

[6] László Babai, Xi Chen, Xiaorui Sun, Shang-Hua Teng, and John Wilmes. 2013. Faster canonical forms for strongly regular graphs. In *IEEE Symposium on Foundations of Computer Science (FOCS)*. 157–166. https://doi.org/10.1109/FOCS.2013.25

[7] László Babai, Paolo Codenotti, Joshua A Grochow, and Youming Qiao. 2011. Code equivalence and group isomorphism. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1395–1408. https://doi.org/10.1137/1.9781611973082.107

[8] László Babai, Paolo Codenotti, and Youming Qiao. 2012. Polynomial-time isomorphism test for groups with no abelian normal subgroups. In *International Colloquium on Automata, Languages, and Programming (ICALP)*. Springer, 51–62. https://doi.org/10.1007/978-3-642-31594-7_5

[9] László Babai, Paul Erdós, and Stanley M Selkow. 1980. Random graph isomorphism. *SIAM Journal on computing* 9, 3 (1980), 628–635. https://doi.org/10.1137/0209047

[10] László Babai and Eugene M Luks. 1983. Canonical labeling of graphs. In *ACM Symposium on Theory of computing (STOC)*. 171–183. https://doi.org/10.1145/800061.808746

[11] László Babai and Youming Qiao. 2012. Polynomial-time isomorphism test for groups with Abelian Sylow towers. In *International Symposium on Theoretical Aspects of Computer Science (STACS)*. 453. https://doi.org/10.4230/LIPIcs.STACS.2012.453

[12] László Babai and John Wilmes. 2013. Quasipolynomial-time canonical form for Steiner designs. In *ACM Symposium on Theory of Computing (STOC)*. https://doi.org/10.1145/2488608.2488642

[13] Reinhold Baer. 1938. Groups with abelian central quotient group. *Trans. Amer. Math. Soc.* 44, 3 (1938), 357–386. https://doi.org/10.1090/S0002-9947-1938-1501972-1

[14] Jendrik Brachter and Pascal Schweitzer. 2020. On the Weisfeiler-Leman dimension of finite groups. In *ACM/IEEE Symposium on Logic in Computer Science (LICS)*. 287–300. https://doi.org/10.1145/3373718.3394786

[15] Peter A Brooksbank, Joshua A Grochow, Yinan Li, Youming Qiao, and James B Wilson. 2019. Incorporating Weisfeiler-Leman into algorithms for group isomorphism. *arXiv preprint arXiv:1905.02518* (2019).

[16] Peter A Brooksbank, Yinan Li, Youming Qiao, and James B Wilson. 2020. Improved algorithms for alternating matrix space isometry: from theory to practice. In *European Symposium on Algorithms (ESA)*. https://doi.org/10.4230/LIPIcs.ESA.2020.26

[17] Peter A Brooksbank, Joshua Maglione, and James B Wilson. 2015. A fast isomorphism test for groups of genus 2. *arXiv preprint arXiv:1508.03033* (2015).

[18] Peter A Brooksbank and James Wilson. 2012. Computing isometry groups of Hermitian maps. *Trans. Amer. Math. Soc.* 364, 4 (2012), 1975–1996. https://doi.org/10.1090/S0002-9947-2011-05388-2

[19] Xi Chen, Xiaorui Sun, and Shang-Hua Teng. 2013. Multi-stage design for quasipolynomial-time isomorphism testing of Steiner 2-systems. In *ACM Symposium on Theory of Computing (STOC)*. https://doi.org/10.1145/2488608.2488643

[20] Samir Datta, Nutan Limaye, Prajakta Nimbhorkar, Thomas Thierauf, and Fabian Wagner. 2009. Planar graph isomorphism is in log-space. In *IEEE Conference on Computational Complexity (CCC)*. 203–214. https://doi.org/10.1109/CCC.2009.16

[21] Heiko Dietrich and James B Wilson. 2022. Group isomorphism is nearly-linear time for most orders. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 457–467. https://doi.org/10.1109/FOCS52979.2021.00053

[22] Volkmar Felsch and Joachim Neubüser. 1970. On a programme for the determination of the automorphism group of a finite group. In *Computational Problems in Abstract Algebra*. 59–60. https://doi.org/10.1016/B978-0-08-012975-4.50011-4

[23] Harley Flanders. 1962. On spaces of linear transformations with bounded rank. *Journal of the London Mathematical Society* 1, 1 (1962), 10–16. https://doi.org/10.1112/jlms/s1-37.1.10

[24] Michael R Garey and David S Johnson. 1979. Computers and intractability: a guide to the theory of NP-completeness.

[25] Joshua A Grochow and Youming Qiao. 2021. On p-group isomorphism: search-to-decision, counting-to-decision, and nilpotency class reductions via tensors. In *36th Computational Complexity Conference (CCC)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik. https://doi.org/10.4230/LIPIcs.CCC.2021.16

[26] Joshua A Grochow and Youming Qiao. 2021. On the complexity of isomorphism problems for tensors, groups, and polynomials I: tensor isomorphism-completeness. In *Innovations in Theoretical Computer Science Conference (ITCS)*. https://doi.org/10.4230/LIPIcs.ITCS.2021.31

[27] Martin Grohe and Daniel Neuen. 2019. Canonisation and definability for graphs of bounded rank width. In *ACM/IEEE Symposium on Logic in Computer Science (LICS)*. 1–13. https://doi.org/10.1109/LICS.2019.8785682

[28] Martin Grohe, Daniel Neuen, and Pascal Schweitzer. 2020. A faster isomorphism test for graphs of small degree. *SIAM J. Comput.* (2020). https://doi.org/10.1137/19M1245293

[29] Martin Grohe, Daniel Wiebking, and Daniel Neuen. 2020. Isomorphism testing for graphs excluding small minors. In *Annual Symposium on Foundations of Computer Science (FOCS)*. 625–636. https://doi.org/10.1109/FOCS46700.2020.00064

[30] Gábor Ivanyos and Youming Qiao. 2019. Algorithms based on*-algebras, and their applications to isomorphism of polynomials with one secret, group isomorphism, and polynomial identity testing. *SIAM J. Comput.* 48, 3 (2019), 926–963. https://doi.org/10.1137/18M1165682

[31] 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 Conference*. Springer, 251–281. https://doi.org/10.1007/978-3-030-36030-6_11

[32] Telikepalli Kavitha. 2007. Linear time algorithms for abelian group isomorphism and related problems. *J. Comput. System Sci.* 73, 6 (2007), 986–996. https://doi.org/10.1016/j.jcss.2007.03.013

[33] Sandra Kiefer, Ilia Ponomarenko, and Pascal Schweitzer. 2019. The Weisfeiler–Leman dimension of planar graphs is at most 3. *Journal of the ACM (JACM)* 66, 6 (2019), 1–31. https://doi.org/10.1145/3333003

[34] François Le Gall. 2009. Efficient isomorphism testing for a class of group extensions. In *International Symposium on Theoretical Aspects of Computer Science (STACS)*. 625–636. https://doi.org/10.4230/LIPIcs.STACS.2009.1830

[35] Mark L Lewis and James B Wilson. 2010. Isomorphism in expanding families of indistinguishable groups. *arXiv preprint arXiv:1010.5466* (2010).

[36] Yinan Li and Youming Qiao. 2017. Linear algebraic analogues of the graph isomorphism problem and the Erdős-Rényi model. In *Annual Symposium on Foundations of Computer Science (FOCS)*. 463–474. https://doi.org/10.1109/FOCS.2017.49

[37] Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. 2017. Fixed-parameter tractable canonization and isomorphism test for graphs of bounded treewidth. *SIAM J. Comput.* 46, 1 (2017), 161–189. https://doi.org/10.1137/140999980

[38] Eugene M Luks. 2015. Group isomorphism with fixed subnormal chains. *arXiv preprint arXiv:1511.00151* (2015).

[39] Gary L. Miller. 1978. On the $n^{\log n}$ isomorphism technique: a preliminary report. In *ACM Symposium on Theory of Computing (STOC)*. 51–58. https://doi.org/10.1145/800133.804331

[40] Daniel Neuen. 2022. Isomorphism testing for graphs excluding small topological subgraphs. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1411–1434. https://doi.org/10.1137/1.9781611977073.59

[41] Eamonn A O'Brien. 1994. Isomorphism testing for p-groups. *Journal of Symbolic Computation* 17, 2 (1994), 133–147. https://doi.org/10.1006/jsco.1994.1007

[42] Youming Qiao, Jayalal Sarma, and Bangsheng Tang. 2012. On isomorphism testing of groups with normal Hall subgroups. *Journal of Computer Science and Technology* 27, 4 (2012), 687–701. https://doi.org/10.1007/s11390-012-1255-7

[43] David J. Rosenbaum. 2013. Bidirectional collision detection and faster deterministic isomorphism testing. *arXiv preprint arXiv:1304.3935* (2013).

[44] David J. Rosenbaum and Fabian Wagner. 2015. Beating the generator-enumeration bound for p-group isomorphism. *Theoretical Computer Science* 593 (2015), 16–25. https://doi.org/10.1016/j.tcs.2015.05.036

[45] Carla Diane Savage. 1980. *An $O(n^2)$ algorithm for abelian group isomorphism*. North Carolina State University.

[46] Tyler Schrock. 2019. *On the complexity of isomorphism in finite group theory and symbolic dynamics*. Ph.D. Dissertation. University of Colorado at Boulder.

[47] Daniel A. Spielman. 1996. Faster isomorphism testing of strongly regular graphs. In *ACM Symposium on Theory of Computing (STOC)*. 576–584. https://doi.org/10.1145/237814.238006

[48] Xiaorui Sun and John Wilmes. 2015. Faster canonical forms for primitive coherent configurations. In *ACM Symposium on Theory of Computing (STOC)*. 693–702. https://doi.org/10.1145/2746539.2746617

[49] Gang Tang, Dung Hoang Duong, Antoine Joux, Thomas Plantard, Youming Qiao, and Willy Susilo. 2022. Practical post-quantum signature schemes from isomorphism problems of trilinear forms. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*. 582–612. https://doi.org/10.1007/978-3-031-07082-2_21

[50] Narayan Vikas. 1996. An $O(n)$ algorithm for abelian p-group isomorphism and an $O(n \log n)$ algorithm for abelian group isomorphism. *J. Comput. System Sci.* 53, 1 (1996), 1–9. https://doi.org/10.1006/jcss.1996.0045

[51] UHM Webb. 1983. On the rank of a p-group of class 2. *Canad. Math. Bull.* 26, 1 (1983), 101–105. https://doi.org/10.4153/CMB-1983-015-5

[52] Boris Weisfeiler and A. A. Lehman. 1968. A reduction of a graph to a canonical form and an algebra arising during this reduction. *Nauchno-Technicheskaya Informatsiya* 9 (1968), 12–16.

[53] Daniel Wiebking. 2020. Graph isomorphism in quasipolynomial time parameterized by treewidth. In *International Colloquium on Automata, Languages, and Programming (ICALP)*. https://doi.org/10.4230/LIPIcs.ICALP.2020.103

[54] James B Wilson. 2009. Decomposing p-groups via Jordan algebras. *Journal of Algebra* 322, 8 (2009), 2642–2679. https://doi.org/10.1016/j.jalgebra.2009.07.029

[55] James B Wilson. 2009. Finding central decompositions of p-groups. *Journal of Group Theory* 12, 6 (2009), 813–830. https://doi.org/10.1515/JGT.2009.015

[56] Viktor N Zemlyachenko, Nickolay M Korneenko, and Regina I Tyshkevich. 1985. Graph isomorphism problem. *Journal of Soviet Mathematics* 29, 4 (1985), 1426–1481. https://doi.org/10.1007/BF02104746

Received 2022-11-07; accepted 2023-02-06