Algorithms based on $*$-algebras, and their applications to isomorphism of polynomials with one secret, group isomorphism, and polynomial identity testing

Gábor Ivanyos* Youming Qiao†

Abstract

We consider two basic algorithmic problems concerning tuples of (skew-)symmetric matrices. The first problem asks to decide, given two tuples of (skew-)symmetric matrices $(B_1, \dots, B_m)$ and $(C_1, \dots, C_m)$, whether there exists an invertible matrix A such that for every $i \in \{1, \dots, m\}$, $A^t B_i A = C_i$. We show that this problem can be solved in randomized polynomial time over finite fields of odd size, the reals, and the complex numbers. The second problem asks to decide, given a tuple of square matrices $(B_1, \dots, B_m)$, whether there exist invertible matrices A and D, such that for every $i \in \{1, \dots, m\}$, AB_iD is (skew-)symmetric. We show that this problem can be solved in deterministic polynomial time over fields of characteristic not 2. For both problems we exploit the structure of the underlying $*$-algebras (algebras with an involutive anti-automorphism), and utilize results and methods from the module isomorphism problem.

Applications of our results range from multivariate cryptography, group isomorphism, to polynomial identity testing. Specifically, these results imply efficient algorithms for the following problems. (1) Test isomorphism of quadratic forms with one secret over a finite field of odd size. This problem belongs to a family of problems that serves as the security basis of certain authentication schemes proposed by Patarin (Eurocrypt 1996). (2) Test isomorphism of p-groups of class 2 and exponent p (p odd) with order p^ℓ in time polynomial in the group order, when the commutator subgroup is of order $p^{O(\sqrt{\ell})}$. (3) Deterministically reveal two families of singularity witnesses caused by the skew-symmetric structure. This represents a natural next step for the polynomial identity testing problem, in the direction set up by the recent resolution of the non-commutative rank problem (Garg-Gurvits-Oliveira-Wigderson, FOCS 2016; Ivanyos-Qiao-Subrahmanyam, ITCS 2017).

1 Introduction

We consider two basic algorithmic problems concerning tuples of (skew-)symmetric matrices. For convenience, for $\epsilon \in \{1, -1\}$, we say an $n \times n$ matrix B is ϵ-symmetric, if $B^t = \epsilon B$. Clearly, when $\epsilon = 1$ (resp. $\epsilon = -1$), B is symmetric (resp. skew-symmetric).

The first problem asks to decide, given two tuples of $n \times n$ ϵ-symmetric matrices $(B_1, \dots, B_m)$ and $(C_1, \dots, C_m)$, whether there exists an invertible $n \times n$ matrix A, such that $\forall i \in [m]$, $A^t B_i A = C_i$. We call this problem *the isometry problem for ϵ-symmetric matrix tuples*. We show that this problem can be solved in randomized polynomial time when the underlying field is a finite field of odd size, the field of real numbers, or the field of complex numbers.

The second problem asks to decide, given a tuple of $n \times n$ matrices $(B_1, \dots, B_m)$, whether there exist invertible $n \times n$ matrices A and D, such that $\forall i \in [m]$, AB_iD is ϵ-symmetric. We call this problem the *ϵ-symmetrization problem for matrix tuples*. We show that this problem can be solved in deterministic polynomial time, as long as the underlying field is not of characteristic 2.

At first sight, these two problems seem to be of interest mostly in computer algebra. However, as we explain below, these results are motivated by, and therefore have applications to, three seemingly unrelated research topics. These are multivariate cryptography, group isomorphism problem, and polynomial identity testing problem, which are traditionally studied in cryptography, computational group theory, and algebraic complexity theory, respectively. The algorithm for isometry testing of ϵ-symmetric matrix tuples leads to substantial improvements over recent algorithms from multivariate cryptography and group isomorphism [BFP15, BMW17]. In particular, the algorithm for isometry testing of symmetric matrix tuples completely settles the so-called Isomorphism of Quadratic Polynomials with One Secret problem over finite fields of odd size [Pat96]. The algorithm for the ϵ-symmetrization problem represents a natural next step for the polynomial identity testing problem in the

*Institute for Computer Science and Control, Hungarian Academy of Sciences, Budapest, Hungary (Gabor.Ivanyos@sztaki.mta.hu).

†Centre for Quantum Software and Information, University of Technology Sydney, Australia (Youming.Qiao@uts.edu.au)

direction set up by the recent resolution of the non-commutative rank problem [GGOW16,IQS17b,IQS17a].

The algorithms for the isometry problem and the ϵ-symmetrization problem share two key ingredients in common. The first one is to utilize the structure of $*$-algebras, that is algebras with an involutive anti-automorphism, underlying these problems. Our use of $*$-algebras is inspired by the works of J. B. Wilson, who pioneered the use of $*$-algebras in computing with p-groups [Wil09a, Wil09b, BW12]. The second one is the results and methods from the module isomorphism problem, which asks to decide, given two tuples of matrices $(B_1, \dots, B_m)$, $(C_1, \dots, C_m)$, whether there exists an invertible matrix A, such that $\forall i \in [m]$, $AB_i = C_iA$. This problem admits two deterministic efficient algorithms by [CIK97,IKS10] and [BL08]. These results and the techniques are used frequently in both algorithms.

This introduction serves as an extended abstract. From Section 1.1 to 1.3, we elaborate on the applications. Since the applications span across three different areas, in order to provide the contexts for readers with different backgrounds, we shall not refrain from including certain background information, despite that it is well-known for researchers in the respective area. In Section 1.4, we formally present the results, explain more on the two key ingredients shared by both algorithms, and describe some open problems. In Section 1.5, we give outlines of the algorithms. The rest of this article then devotes to detailed descriptions of the algorithms.

We now set up some notation. $\mathbb{F}$, $\mathbb{E}$, and $\mathbb{K}$ are used to denote fields. $\mathbb{F}_q$ denotes the finite field of size q, $\mathbb{R}$ the real field, and $\mathbb{C}$ the complex field. Unless otherwise stated, we work with fields of characteristic not 2. $M(n, \mathbb{F})$ denotes the linear space of $n \times n$ matrices over $\mathbb{F}$, and $\mathrm{GL}(n, \mathbb{F})$ the group of invertible matrices in $M(n, \mathbb{F})$. $S^\epsilon(n, \mathbb{F})$ denotes the linear space of $n \times n$ ϵ-symmetric matrices over $\mathbb{F}$. We may write $M(n, q)$, $\mathrm{GL}(n, q)$, and $S^\epsilon(n, q)$ for $M(n, \mathbb{F}_q)$, $\mathrm{GL}(n, \mathbb{F}_q)$, and $S^\epsilon(n, \mathbb{F}_q)$, respectively. A *matrix space* is a linear subspace of $M(n, \mathbb{F})$, and $\langle\cdot\rangle$ denotes linear span. Let $\mathbf{B} = (B_1, \dots, B_m) \in M(n, \mathbb{F})^m$ be a matrix tuple. For $A, D \in M(n, \mathbb{F})$, $A\mathbf{B}D := (AB_1D, \dots, AB_mD)$ and $\mathbf{B}^t := (B_1^t, \dots, B_n^t)$.

1.1 Multivariate cryptography In 1996, Patarin proposed a family of asymmetric cryptography schemes based on equivalence of polynomials in [Pat96], which can be used for authentication and signature. One scheme in this family is based on the assumed hardness of the following problem.

Problem 1.1. (Isomorphism of Quadratic Forms with One Secret (IQF1S)) *Let $\mathbf{f} = (f_1, \dots, f_m)$ and $\mathbf{g} = (g_1, \dots, g_m)$ be two tuples of homogeneous quadratic polynomials in n variables $\{x_1, \dots, x_n\}$ over a finite field $\mathbb{F}$. Decide if there exists $A \in \mathrm{GL}(n, \mathbb{F})$ such that $\forall k \in [m]$, $f_k^A = g_k$, where $A = (a_{i,j})_{i,j\in[n]}$ acts on $\{x_1, \dots, x_n\}$ by sending x_i to $\sum_{j\in[n]} a_{i,j}x_j$.*

For readers familiar with Patarin's work [Pat96], IQF1S is Patarin's Isomorphism of Polynomials with One Secret (IP1S) restricting to quadratic polynomials, which asks the same question but for possibly inhomogeneous quadratic polynomials and affine transformations.[1] Such a restriction is well justified from the practical viewpoint, as it minimizes the public-key storage and improves the actual performance, so this has been studied most in the literature. Since Patarin's introduction of these problems, IQF1S and several related problems have been intensively studied [PGC98, GMS03, Per05, FP06, Kay11, BFFP11, MPG13, BFV13, PFM14, BFP15].

Most notably, in [BFP15], Berthomieu et al. presented an efficient randomized algorithm for IQF1S under the conditions that (1) $\mathbf{f}$ satisfies a regularity condition, namely that there exists a nondegenerate form in the linear span of f_i's, (2) the underlying field is large enough and of characteristic not 2, and (3) the desired solution may be from an extension field [BFP15, Theorem 2]. They further observed that, it seems that most known algorithms on IQF1S would fail on the irregular instances, and proposed the complexity of such instances as an open question [BFP15, Sec. 1, Open Question].

By the classical correspondence between quadratic forms and symmetric matrices, it is easy to see the equivalence between IQF1S and the isometry problem of tuples of symmetric matrices. Our algorithm for the latter problem then translates to a complete solution of IQF1S over finite fields of odd size, answering [BFP15, Sec. 1, Open Question] for such fields.

Theorem 1.1. *Let $\mathbb{F}$ be a finite field of odd size. There exists a randomized polynomial-time algorithm that solves the Isomorphism of Quadratic Forms with One Secret problem over $\mathbb{F}$.*

Furthermore, there has been a large body of works which aim to build public key cryptography schemes based on the hardness of solving systems of quadratic polynomials over finite fields. This approach is regarded as one candidate for post-quantum cryptography, in particular as a signature scheme [CJL+16]. We refer the reader to the thesis of Wolf [Wol05] for an overview, and the recent article [PCDY17] and references therein

[1] Patarin's formulation is known to reduce to the formulation here [BFP15, Proposition 5].

for recent advances in this area. IQF1S and related problems play an important role in such schemes. As pointed out in [Wol05, Sec. 2.6.1], though often not explicitly stated, it seems crucial to assume that IQF1S and related problems are difficult to ensure the security of these schemes. Theorem 1.1 then suggests that the "one-secret" versions of such schemes based on quadratic polynomials may not be secure.

1.2 Group isomorphism problem Group isomorphism problem (GpI) asks to decide whether two finite groups of order n are isomorphic. It has been studied for several decades in both Computational Group Theory (CGT) and Theoretical Computer Science. The difficulty of this problem depends crucially on how we represent the groups in the algorithms. If the goal is to obtain an algorithm running in time poly(n), then we may assume that we have at our disposal the Cayley (multiplication) table of the group, as the Cayley table can be recovered from most reasonable models for computing with finite groups in time poly(n). Therefore, we restrict our discussion mostly to this very redundant model, which is meaningful mainly because we do not know a poly(n)-time or even an $n^{o(\log n)}$-time algorithm [Wil14] (log to the base 2), despite that a simple $n^{\log n+O(1)}$-time algorithm has been known for decades [FN70, Mil78]. The past few years have witnessed a resurgence of activity on algorithms for this problem with worst-case analyses in terms of the group order; we refer the reader to [GQ17] which contains a survey of these algorithms.

It is long believed that p-groups (groups of a prime power order) form the bottleneck case for GpI. In fact, the decades-old quest for a polynomial-time algorithm has focused on class-2 p-groups, with little success. Even if we restrict further to p-groups of class 2 and exponent p, the problem is still difficult. Recently, some impressive progress on such p-groups was made on the CGT side, as seen in the works of Wilson, Brooksbank, and their collaborators [Wil09a, LW12, BMW17].

Most notably, a main result in [BMW17] is a polynomial-time algorithm for p-groups of class 2 and exponent p, when the commutator subgroup is of order p^2, in the model of quotients of permutation groups [KL90]. This of course settles the same case in the Cayley table model. In fact, the same class of groups in the Cayley table model can be handled using one specific technique called the Pfaffian isomorphism test in [BMW17, Sec. 6.2]. Still, despite all the progress, an efficient algorithm for p-groups of class 2 and exponent p, with the commutator subgroup of order even p^3, was not known in the Cayley table model. Since we now have an efficient algorithm to test isometry of tuples of skew-symmetric matrices, the following result can be established.

THEOREM 1.2. *Let p be an odd prime, and let two p-groups of class 2 and exponent p of order p^ℓ, G and H, be given by Cayley tables. If the commutator subgroup of G is of order $p^{O(\sqrt{\ell})}$, then there exists a deterministic polynomial-time algorithm to test whether G and H are isomorphic.*

We explain how to obtain Theorem 1.2 from our result. While the following reduction is well-known in CGT, we include it here for readers from other areas. Given a class 2 and exponent p p-group G, let $[G,G]$ denote its commutator subgroup. Due to the exponent p and class 2 condition, we have $G/[G,G] \cong \mathbb{Z}_p^n$ and $[G,G] \cong \mathbb{Z}_p^m$ for some n and m such that $n+m=\ell$. Fixing bases of $G/[G,G]$ and $[G,G]$, and taking the commutator bracket, we obtain a skew-symmetric bilinear map $b_G : \mathbb{F}_p^n \times \mathbb{F}_p^n \to \mathbb{F}_p^m$, represented by $\mathbf{B} \in S^{-1}(n,p)^m$. For H to be isomorphic to G, it is necessary that $\dim_{\mathbb{Z}_p}(H/[H,H]) = \dim_{\mathbb{Z}_p}(G/[G,G])$ and $\dim_{\mathbb{Z}_p}([H,H]) = \dim_{\mathbb{Z}_p}([G,G])$, so by the same construction we obtain another $\mathbf{C} \in S^{-1}(n,p)^m$. We then need the following definition.

DEFINITION 1.1. *Given $\mathbf{B} = (B_1, \dots, B_m)$ and $\mathbf{C} = (C_1, \dots, C_m)$ from $S^\epsilon(n,\mathbb{F})$, $\mathbf{B}$ and $\mathbf{C}$ are* pseudo-isometric, *if there exists $X \in \mathrm{GL}(n,\mathbb{F})$ such that $\langle X^t B_1 X, \dots, X^t B_m X\rangle = \langle C_1, \dots, C_m\rangle$.*

The key connection then is Baer's correspondence, which, put in this context, gives that G and H are isomorphic if and only if $\mathbf{B}$ and $\mathbf{C}$ are pseudo-isometric [Bae38]. By the condition that $m = O(\sqrt{\ell})$, we can enumerate all bases of $\mathbf{C}$ at a multiplicative cost of $p^{m^2} = p^{O(\ell)}$, and for each fixed basis, apply the algorithm for isometry testing. This gives Theorem 1.2.

As Brooksbank and Wilson have communicated to us, our algorithm may be useful in some models studied in CGT. Also, in multivariate cryptography, the problem Isomorphism of Quadratic Forms with Two Secrets (IQF2S) just asks to test the pseudo-isometry of tuples of symmetric matrices. Formally, the IQF2S problem asks to decide, given $\mathbf{B}, \mathbf{C} \in S^1(n,\mathbb{F})$, whether they are pseudo-isometric. Therefore a result analogous to Theorem 1.2 can be obtained for IQF2S.

1.3 Polynomial identity testing Fix $\epsilon \in \{1,-1\}$. Let us see how to cast the ϵ-symmetrization problem as an instance of the polynomial identity testing problem. Given $\mathbf{B} = (B_1, \dots, B_m) \in M(n,\mathbb{F})^m$, there exist invertible matrices A, D such that $\forall i \in [m]$, AB_iD is ϵ-symmetric if and only if $\forall i \in [m]$, $D^{-t}AB_i = D^{-t}(AB_iD)D^{-1}$ is ϵ-symmetric. Therefore we can reduce to finding an invertible matrix E such that $\forall i \in [m]$, EB_i is ϵ-symmetric. Suppose for now that E is a matrix of variables. The equations $\forall i \in [m], EB_i = \epsilon B_i^t E^t$ set up a system of linear forms in these variables. Let $C_1, \dots, C_\ell$ be a linear basis of the solution space, and $\mathcal{C}$ be the matrix space $\langle C_1, \dots, C_\ell \rangle \leq M(n, \mathbb{F})$. The problem then becomes to decide whether $\mathcal{C}$ contains an invertible matrix. To decide whether a matrix space, given by a linear basis, contains only non-invertible matrix is known as the symbolic determinant identity testing (SDIT) problem, which is equivalent to the polynomial identity testing (PIT) for weakly skew arithmetic circuits [Tod92][2].

When $|\mathbb{F}| = \Omega(n)$, SDIT admits a randomized efficient algorithm via the Schwartz-Zippel lemma. To devise a deterministic efficient algorithm for SDIT is a major problem in algebraic complexity theory due to its implication to arithmetic circuit lower bounds. Specifically, in [CIKK15] (building on [KI04]), Carmosino et al. show that such an algorithm implies the existence of a polynomial family such that its graph is in NE, but it cannot be computed by polynomial-size arithmetic circuits. Such a lower bound is generally considered to be beyond current techniques, and would be recognized as a breakthrough if established. The research into PIT has received quite attention since early 2000's (see the surveys [Sax09, SY10, Sax13]).

Our algorithm for the ϵ-symmetrization problem then provides a deterministic solution to this specific instance of SDIT. Our motivation to look at this problem at the first place was from the recent resolution of the non-commutative rank problem by Garg et al. [GGOW16] and Ivanyos et al. [IQS17b, IQS17a], and the intricate relation between the non-commutative rank problem and SDIT, which we explain below.

A matrix space $\mathcal{B} \leq M(n, \mathbb{F})$ is non-singular, if $\mathcal{B}$ contains an invertible matrix, and singular otherwise. SDIT then asks to decide whether a matrix space is singular. To obtain an arithmetic circuit lower bound via [CIKK15], it is actually enough to put SDIT in NP, that is, to find a short witness that helps to testify the singularity of singular matrix spaces. One such singularity witness, which is the reminiscent of the "shrunk subset" as in Hall's marriage theorem for bipartite graphs, and closely related to the linear matroid intersection problem [Lov89], is the following. For $\mathcal{B} \leq M(n, \mathbb{F})$, $U \leq \mathbb{F}^n$ is a shrunk subspace of $\mathcal{B}$, if $\dim(U) > \dim(\mathcal{B}(U))$ where $\mathcal{B}(U) = \langle B(U) : B \in \mathcal{B} \rangle$. The decision version of the non-commutative rank problem then asks to decide whether $\mathcal{B}$ has a shrunk subspace. Deterministic efficient algorithms for the non-commutative rank problem were recently devised in [GGOW16] (over $\mathbb{Q}$) and in [IQS17b, IQS17a] (over any field).

A direct consequence of settling the non-commutative rank problem on SDIT is that we can restrict our attention to those singular matrix spaces without a shrunk subspace, which we call exceptional spaces. As described by Lovász in [Lov89] (see also [Atk83, EH88]), the skew-symmetric structure naturally yields two families of exceptional spaces. To introduce them we need the following definition. Two matrix spaces $\mathcal{B}, \mathcal{C} \leq M(n, \mathbb{F})$ are *equivalent*, if there exist $A, D \in \mathrm{GL}(n, \mathbb{F})$ such that $A\mathcal{B}D = \mathcal{C}$ (equal as subspaces). Note that whether a matrix space is singular is preserved by the equivalence relation. We now list the two families from [Lov89].

(1) If n is odd and $\mathcal{B} \leq M(n, \mathbb{F})$ is equivalent to a subspace in $S^{-1}(n, \mathbb{F})$, then $\mathcal{B}$ is singular, as every skew-symmetric matrix is of even rank.

(2) Given $C_1, \dots, C_n \in S^{-1}(n, \mathbb{F})$, let $\mathcal{C} \leq M(n, \mathbb{F})$ consist of all the matrices of the form $[C_1 v, C_2 v, \dots, C_n v]$ over $v \in \mathbb{F}^n$. Since $v^t[C_1 v, C_2 v, \dots, C_n v] = [v^t C_1 v, v^t C_2 v, \dots, v^t C_n v] = 0$, $\mathcal{C}$ is singular, and we call such $\mathcal{C}$ a skew-symmetric induced matrix space. If $\mathcal{B}$ is equivalent to a skew-symmetric induced matrix space, then $\mathcal{B}$ is singular as well. Note that w.l.o.g. we can assume that $\mathcal{B}$ is a subspace of $M(n, \mathbb{F})$ of dimension n.

These two families of exceptional matrix spaces can be deterministically recognized as follows.

THEOREM 1.3. *Let $\mathbb{F}$ be a field of characteristic not 2. Given $\mathcal{B} = \langle B_1, \dots, B_m \rangle \leq M(n, \mathbb{F})^m$, there exists a deterministic polynomial-time algorithm that decides whether $\mathcal{B}$ is equivalent to a subspace in $S^{-1}(n, \mathbb{F})$, or a skew-symmetric induced matrix space.*

We explain how Theorem 1.3 follows from our ϵ-symmetrization algorithm. The case (1) is straightforward: apply the skew-symmetrization algorithm to the given linear basis of $\mathcal{B}$. In case (2), suppose $B_i = [b_{i,1}, \dots, b_{i,n}]$ where $b_{i,j} \in \mathbb{F}^n$, $j \in [n]$ are the columns of B_i. Following an observation of Lovász in [Lov89], construct $B_i' = [b_{1,i}, \dots, b_{n,i}]$ for $i \in [n]$. It can be verified that $\mathcal{B}$ is equivalent to some $\mathcal{C}$ of the form described in (2) if and only if $\mathcal{B}' = \langle B_1', \dots, B_n' \rangle$ is equivalent to a subspace in $S^{-1}(n, \mathbb{F})$. We can then apply the skew-symmetrization algorithm to $(B_1', \dots, B_n')$ to conclude.

[2] An arithmetic circuit is weakly skew if each product gate is of fan-in 2 and has at least one child such that the subcircuit rooted at it is separate from the other parts of the circuit [Tod92, MP08]. The computation power of weakly skew circuit is known to be equivalent to the model of symbolic determinants, and between arithmetic formulas and arithmetic circuits.

1.4 Results and techniques

Statement of the results. We first define three equivalence relations for matrix tuples.

DEFINITION 1.2. *Let* $\mathbf{B} = (B_1, \dots, B_m), \mathbf{C} = (C_1, \dots, C_m) \in M(n, \mathbb{F})^m$. $\mathbf{B}$ *and* $\mathbf{C}$ *are* conjugate*, if* $\exists A \in \mathrm{GL}(n, \mathbb{F})$*, such that* $A\mathbf{B} = \mathbf{C}A$. *They are* equivalent*, if* $\exists A, D \in \mathrm{GL}(n, \mathbb{F})$*, such that* $A\mathbf{B} = \mathbf{C}D$. *They are* isometric*, denoted as* $\mathbf{B} \sim \mathbf{C}$*, if* $\exists A \in \mathrm{GL}(n, \mathbb{F})$*, such that* $A^t\mathbf{B}A = \mathbf{C}$*; such an* A *is called an isometry from* $\mathbf{B}$ *to* $\mathbf{C}$.

We show that testing whether two ϵ-symmetric matrix tuples are isometric can be solved efficiently over $\mathbb{F}_q$ with q odd, $\mathbb{R}$, and $\mathbb{C}$. Note that the algorithm for $\mathbb{F}_q$ is probabilistic.

THEOREM 1.4. *1. (Finite fields of odd size) Given* $\mathbf{B}, \mathbf{C} \in S^\epsilon(n, q)^m$ *with* q *odd, there exists a randomized polynomial-time algorithm that decides whether* $\mathbf{B}$ *and* $\mathbf{C}$ *are isometric. If* $\mathbf{B}$ *and* $\mathbf{C}$ *are isometric, the algorithm also computes an explicit isometry in* $\mathrm{GL}(n, q)$. *This algorithm can be derandomized at the price of running in time* $\mathrm{poly}(n, m, \log q, p)$ *where* $p = \mathrm{char}(\mathbb{F}_q)$.

2. (The real field $\mathbb{R}$*) Let* $\mathbb{E} \subseteq \mathbb{R}$ *be a number field. Given* $\mathbf{B}, \mathbf{C} \in S^\epsilon(n, \mathbb{E})^m$*, there exists a deterministic polynomial-time algorithm that decides whether* $\mathbf{B}$ *and* $\mathbf{C}$ *are isometric over some number field* $\mathbb{K}$ *such that* $\mathbb{E} \subseteq \mathbb{K} \subseteq \mathbb{R}$. *If* $\mathbf{B}$ *and* $\mathbf{C}$ *are indeed isometric, the algorithm also computes an explicit isometry, represented as a product of matrices, where each matrix is over some extension field of* $\mathbb{E}$ *of extension degree* $\mathrm{poly}(n, m)$.

3. (The complex field $\mathbb{C}$*) Let* $\mathbb{E}$ *be a number field. Given* $\mathbf{B}, \mathbf{C} \in S^\epsilon(n, \mathbb{E})^m$*, there exists a deterministic polynomial-time algorithm that decides whether* $\mathbf{B}$ *and* $\mathbf{C}$ *are isometric over some number field* $\mathbb{K}$ *such that* $\mathbb{E} \subseteq \mathbb{K}$. *If* $\mathbf{B}$ *and* $\mathbf{C}$ *are indeed isometric, the algorithm also computes an explicit isometry, represented as a product of matrices, where each matrix is over some extension field of* $\mathbb{E}$ *of extension degree* $\mathrm{poly}(n, m)$.

We call $\mathbf{B} \in M(n, \mathbb{F})^m$ *ϵ-symmetrizable*, if $\mathbf{B}$ is equivalent to a tuple of ϵ-symmetric matrices. Our second main result concerns the problem of testing whether a matrix tuple is ϵ-symmetrizable.

THEOREM 1.5. *Let* $\mathbb{F}$ *be a field of characteristic not 2. Given* $\mathbf{B} \in M(n, \mathbb{F})^m$*, there exists a deterministic algorithm that decides whether* $\mathbf{B}$ *is* ϵ*-symmetrizable, and if it is, computes* $A, D \in \mathrm{GL}(n, \mathbb{F})$ *such that* $A\mathbf{B}D \in S^\epsilon(n, \mathbb{F})^m$. *The algorithm uses polynomially many arithmetic operations. Over a number field the final data as well as all the intermediate data have size polynomial in the input data size, hence the algorithm runs in polynomial time.*

Two key ingredients. Let us first review the concept of $*$-algebras, and see how to get a $*$-algebra from a tuple of ϵ-symmetric matrices. Recall that, a $*$-algebra A is an algebra with $* : A \to A$ being an anti-automorphism of order at most 2. $*$-algebras have been studied since 1930's [Alb39] (see [Lew06] for a recent survey). Let $M(n, \mathbb{F})^{op}$ be the opposite full matrix algebra, which is the ring consisting of all matrices in $M(n, \mathbb{F})$ with the multiplication $\circ$ as $A \circ B = BA$. $*$-algebras arise from ϵ-symmetric matrix tuples by considering the *adjoint algebra* of $\mathbf{B} \in S^\epsilon(n, \mathbb{F})^m$, which consists of $\{(A, D) \in M(n, \mathbb{F})^{op} \oplus M(n, \mathbb{F}) | A^t\mathbf{B} = \mathbf{B}D\}$, with a natural involution $*$ as $(A, D)^* = (D, A)$.

We then turn to the module isomorphism problem (MI). Given $\mathbf{B}, \mathbf{C} \in M(n, \mathbb{F})^m$, MI asks if $\mathbf{B}$ and $\mathbf{C}$ are conjugate. This problem is termed as module isomorphism, as the matrix tuple $\mathbf{B} = (B_1, \dots, B_m)$ can be viewed as a linear representation of a finitely generated algebra generated by m elements. Two deterministic polynomial-time algorithms for MI have been devised in [CIK97, IKS10] and [BL08]. Note that MI may also be cast as an instance of the polynomial identity testing problem like the ϵ-symmetrization problem.

More comparison with previous works. Some comparisons with previous works were already stated in Section 1.1 and 1.2. We now add some more details on the technical side. In Section 1.1, we mentioned the work of Berthomieu et al. [BFP15] which solves the IQF1S possibly over an extension field, for regular instances and large enough fields. Here we seek "rational" solutions (i. e. those over the given base field) in the finte case and seek soltuions over a real extension field. An interesting observation is that the algorithm of Berthomieu et al. may be cast as working with a $*$-algebra, but in a much restricted setting. We explain this in detail in [IQ17, Appendix]. In Section 1.2, we described how our result, when applied to p-group isomorphism, compares to the result of Brooksbank et al. [BMW17]. The relevant technique there, called the Pfaffian isomorphism test [BMW17, Sec. 6.2], is completely different from ours, and seems quite restricted to pairs of skew-symmetric matrices.

The work [BW12] by Brooksbank and Wilson is the most important precursor to our Theorem 1.4. In [BW12], the main result, rephrased in our setting, is an efficient algorithm that, given $\mathbf{B} \in S^\epsilon(n, q)^m$ with q odd, computes a generating set for the group $\{X \in \mathrm{GL}(n, q) \mid X^t\mathbf{B}X = \mathbf{B}\}$. This is exactly the "automorphism version" of the isometry problem.

However, unlike many other isomorphism problems, the isometry problem is not known to reduce to this automorphism version. This is similar to the module isomorphism problem: the automorphism version of MI asks to compute a generating set of the unit group in a matrix algebra, which was solved in [BO08]. The ideas and the techniques for the unit group computation in [BO08] and for MI in [CIK97,IKS10,BL08] are totally different. So Theorem 1.4 cannot be easily deduced as a corollary from [BW12].

Generalizations of the main results. Theorem 1.4 can be generalized to the following setting. Following [BW12], for an linear automorphism $\theta \in \mathrm{GL}(W)$ we call a bilinear map over a field $\mathbb{F}$, $b : V \times V \to W$ *θ-Hermitian*, if for all $u, v \in V$, $b(u, v) = \theta(b(v, u))$. Obviously, nontrivial Hermitian maps exist only if θ^2 is the identity. Hermitian bilinear maps subsume symmetric bilinear maps (θ being the identity matrix) and skew-symmetric bilinear maps (θ being -1 times the identity matrix). It allows for (after fixing bases of V and W) a tuple of mixed symmetric and skew-symmetric matrices. In fact, by a change of basis of W, we may always assume that θ is a diagonal matrix with 1 and -1's on the diagonal and in our arguments and algorithms we only need the replace ϵ by a tuple $(\epsilon_1, \dots, \epsilon_m)$ and equations of type $B_i^t = \epsilon B_i$ by $B_i^t = \epsilon_i B_i$. Furthermore, the concept captures Hermitian forms by [BW12, Sec. 3.1]: for a Hermitian form $b : V \times V \to \mathbb{F}_{q^2}$ where $V \cong \mathbb{F}_{q^2}^n$, we can represent it as a pair of bilinear forms over $\mathbb{F}_q$, $b_1, b_2 : V' \times V' \to \mathbb{F}_q$ where $V' \cong \mathbb{F}_q^{2n}$, and $\theta \in \mathrm{GL}(2, q)$ corresponds to the field involution $\alpha \to \alpha^q$ for $\alpha \in \mathbb{F}_{q^2}$. Hermitian complex or quaternionic matrices are also included: assume that D is a finite dimensional divison algebra over $\mathbb{F}$ with involution $\bar{\cdot} : D \to D$, such that $\mathbb{F}$ coincides with the subfield of the center of D consisting of the elements fixed by $\bar{\cdot}$. Then the map $*$ sending a matrix to the transpose of its elementwise $\bar{\cdot}$-conjugate is an involution on $M(n, D)$, and the matrices invariant under $*$ are called $*$-Hermitian. Indeed, let d be the dimension of D over $\mathbb{F}$. Then we can interpret D and D^n as vector spaces of dimension d resp. dn over $\mathbb{F}$, and a matrix in $M(n, D)$ as an $\mathbb{F}$-bilinear map from $D^n \times D^n$ to D. Then $*$-Hermitian matrices are interpreted as Hermitian bilinar maps for $\bar{\cdot}$. (Naturally, an m-tuple of $*$-Hermitian matrices become a Hermitian map from $D^n \times D^n$ to D^m.)

Interestingly, Theorem 1.4 allows us to solve the isometry problem for a tuple of arbitrary matrices. Given $\mathbf{B}, \mathbf{C} \in M(n, \mathbb{F})^m$, we can construct $\mathbf{B}' = (\frac{1}{2}(B_1 + B_1^t), \dots, \frac{1}{2}(B_m + B_m^t), \frac{1}{2}(B_1 - B_1^t), \dots, \frac{1}{2}(B_1 - B_1^t))$, and similarly $\mathbf{C}'$. Then it is easy to verify that $\mathbf{B} \sim \mathbf{C}$ if and only if $\mathbf{B}' \sim \mathbf{C}'$. Combining with the observation from the last paragraph, we have the following.

Corollary 1.1. *The statement of Theorem 1.4 holds for $\mathbf{B}, \mathbf{C} \in M(n, \mathbb{F}_q)^m$, $M(n, \mathbb{E})^m$ with a number field $\mathbb{E} \subseteq \mathbb{R}$, or $M(n, \mathbb{E})^m$ with a number field $\mathbb{E}$.*

Theorem 1.5 can also be generalized to transforming bilinear maps to θ-Hermitian ones, including the case of tuples of complex and quaternionic matrices.

Some open problems. There are two immediate open problems left.

The first one is to extend both of our results to fields of characteristic 2. While presenting the algorithm for the isometry problem in Section 3, we indicate explicitly in each step whether the characteristic not 2 is required, and one may want to examine those steps where the characteristic not 2 condition is crucial. For the ϵ-symmetrization problem, one may want to start with examining the key lemma, Lemma 1.2, in the setting of characteristic-2 fields.

The second one is to solve the isometry test problem over a number field without going to extension fields. To extend our current approach to deal with the second problem involves certain number-theoretic obstacles even over $\mathbb{Q}$. Namely, our present method relies on representing a simple algebra explicitly as a full matrix algebra over a division ring, but there is a randomized reduction from factoring squarefree integers to this task for a central simple algebra of dimension 4 over $\mathbb{Q}$ assuming the Generalized Riemann Hypothesis [Rón87]. Even deciding whether a four dimensional noncommutative simple algebra over $\mathbb{Q}$ is isomorphic to $M(2, \mathbb{Q})$ is equivalent to deciding deciding quadratic residuosity modulo composite numbers. This kind of obstacles appears to be inherent: a ternary quadratic form over $\mathbb{Q}$ is isotropic if and only if an associated noncommutative simple algebra of dimension four over $\mathbb{Q}$ is isomorphic to $M(2, \mathbb{Q})$. Now consider an indefinite symmetric 3 by 3 matrix B with rational entries having determinant d. Then the ternary quadratic form with Gram matrix B is either anisotropic or isometric to the form having matrix

$$\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & -d \end{pmatrix}.$$

Thus over $\mathbb{Q}$, the isometry problem a single ternary quadratic form is at least as hard as deciding whether an algebra is isomorphic to $M(2, \mathbb{Q})$. Actually, there is a randomized polynomial time reduction from testing whether a simple algebra over a number field $\mathbb{F}$ is isomorphic with a full matrix algebra over $\mathbb{F}$ to factoring integers, see [Rón92] and [IR93] However, for the constructive version of isomorphisms with full matrix algebras such a reduction is only known for the case $M(n, K)$ where n is bounded by a constant, and K is from a finite collection of number fields [IRS12]. Therefore, to determine the relation between the complexity of the isometry problem and that of factoring, it might be useful to devise an alternative approach which gets around constructing explicit isomorphims with full matrix algebras.

Future directions. Given Theorem 1.4, the next target is of course to study IQF2S and isomorphism testing of p-groups of class 2 and exponent p. For these two problems, the first goal would be to design, for $\mathbf{B} \in S^{\epsilon}(n,q)^m$, an algorithm in time $q^{O(n+m)}$. In the context of p-groups of class 2 and exponent p, this amounts to solve isomorphism testing for this group class in time polynomial in the group order, which seems a difficult problem already. By Theorem 1.4, this target seems most difficult when m and n are comparable, say $m = n$. One idea may be to reduce to the parameters m' and n' such that $m' = O(n^{1/2})$ and $n' = \mathrm{poly}(n)$, so that we can use Theorem 1.4 to get an algorithm in time $q^{O(n)}$. It is also noteworthy that recently, Yinan Li and the second author devised an algorithm for $m = \Theta(n)$ in *average-case* time $q^{O(n)}$ [LQ17]; the average-case analysis is done in a random model for linear spaces of skew-symmetric matrices over finite fields, that can be viewed as a linear algebraic analogue of the Erdős-Rényi model for random graphs.

Theorem 1.3 represents a natural step in the direction for derandomizing SDIT set up by the resolution of the non-commutative rank problem [GGOW16, IQS17b, IQS17a]. While most research activities on PIT and SDIT put constraints on the structural properties of the arithmetic circuits [Sax09, SY10, Sax13], this direction puts constraints on the singularity witnesses which are inspired by geometric considerations [EH88] and/or combinatorial considerations [Lov89]. At present, we are not aware of an explicit connection between these two different styles of constraints. It is an interesting question as to whether these geometric and/or combinatorial considerations can be made more systematic to yield a formal strategy to attack SDIT.

1.5 Algorithm outlines We now outline the algorithms, in the hope to illustrate the roles of $*$-algebras and the module isomorphism problem. It should be noted that we have to omit several salient details, and the interested reader is referred to Section 3 and 4 for complete descriptions.

An outline of the main algorithm for Theorem 1.4. Let $\mathbb{F}$ be a field. Recall that we have $\mathbf{B} = (B_1, \ldots, B_m)$ and $\mathbf{C} = (C_1, \ldots, C_m) \in S^{\epsilon}(n, \mathbb{F})^m$. The goal is to decide if there exists $F \in \mathrm{GL}(n, \mathbb{F})$ such that $\forall i \in [m]$, $F^t B_i F = C_i$. The main steps of the algorithm are as follows.

1. *Reduce to the non-degenerate case.* If $\mathbf{B}$ is degenerate, that is $\cap_{i\in[m]} \ker(B_i) \neq \mathbf{0}$, we can reduce to the non-degenerate case by restricting to the non-degenerate part. See Section 3.1.

2. *Solve the twisted equivalence problem.* In this step we test whether $\mathbf{B}$ and $\mathbf{C}$ are "twisted equivalent", that is, whether there exist $A, D \in \mathrm{GL}(n, q)$ such that $A^t\mathbf{B} = \mathbf{C}D$. This problem can be solved efficiently by reducing to the module isomorphism problem. See Section 3.2.

3. *Reduce to decomposing a symmetric element in a $*$-algebra.* At the beginning of this step we know that $\mathbf{B}$ and $\mathbf{C}$ are twisted equivalent under some $A, D \in \mathrm{GL}(n, q)$. Note that if $D = A^{-1}$ then we are done. If not, the hope is to transform A and D appropriately to get an invertible matrix F such that $\mathbf{B}$ and $\mathbf{C}$ are twisted equivalent under F and F^{-1}, if such an F exists. Let $E = A^{-1}D^{-1}$, and define the adjoint algebra of $\mathbf{C}$, $\mathfrak{A} = \mathrm{Adj}(\mathbf{C}) := \{(A, D) \in M(n, \mathbb{F})^{op} \oplus M(n, \mathbb{F}) : \forall i \in [m], A^t C_i = C_i D\}$. It can be verified that $E \in \mathfrak{A}$, and $E^* = E$. The important observation then is that, there exists such F if and only if there exists $X \in \mathfrak{A}$ such that $E = X^* X$. See Section 3.3.

4. *Solve the $*$-symmetric decomposition problem.* This is the main technical piece of this algorithm. This step relies on certain results about the structure of $*$-algebras, which is summarized in Section 2. The basic idea is to utilize the algebra structure of $\mathfrak{A}$, to reduce to the semisimple case, and then further to the simple case. To deal with the simple case turns out to be exactly the isometry problem for a *single* (symmetric, skew-symmetric, or Hermitian...) form, which can be solved using existing algorithms. We now outline the main steps.

   4.a. *Compute the algebra structure of* $\mathfrak{A}$. We start with computing the algebra structure of $\mathfrak{A}$, including the Jacobson radical $J(\mathfrak{A})$, the decomposition of the semisimple quotient into simple summands, and for each simple summand, an explicit isomorphism with a matrix ring over a division algebra. This can be achieved by resorting to known algorithms by Rónyai [Rón90] and Eberly [Ebe91a, Ebe91b]. This step is the main bottleneck to extend this algorithm to number fields (without going to extension fields). See Section 3.4.1.

   4.b. *Recognize the $*$-algebra structure.* We then take into account the $*$-algebra structure. The involution $*$ preserves the Jacobson radical, so it induces an involution on the semisimple quotient, denoted again by $*$. For a particular summand S of the semisimple quotient, $*$ either switches S with another summand, or preserves it. In the the latter case, by the structure theory of $*$-algebras in the simple case, $*$ has to be in a particular form, and this form can be computed explicitly by resorting to the module isomorphism problem. See Section 3.4.2.

4.c. *Reduce to the semisimple case.* In this step, we show that any solution to the $*$-symmetric decomposition problem for $\mathfrak{A}/J(\mathfrak{A})$ and $E + J(\mathfrak{A})$ can be lifted efficiently to a solution to the $*$-symmetric decomposition problem for $\mathfrak{A}$ and E. This procedure crucially relies on that we work with fields of characteristic not 2, and is the main bottleneck to extend this algorithm to fields of characteristic 2. This means that we can reduce to work with semisimple $*$-algebra $\mathfrak{A}$ in the following. See Section 3.4.3.

4.d. *Reduce to the $*$-simple and simple case.* In this step, we want to tackle the $*$-symmetric decomposition problem for a semisimple $*$-algebra $\mathfrak{A}$. Recall that a decomposition of $\mathfrak{A}$ as a sum of simple summands has been computed in Step (4.a). We present a reduction to the same problem for those simple summands that are preserved by $*$. This means that we can reduce to work with a simple $*$-algebra $\mathfrak{A}$. See Section 3.4.4.

4.e. *Tackle the simple case by reducing to the isometry problem for a single form.* In this step, we want to solve the $*$-symmetric decomposition problem for a simple $*$-algebra $\mathfrak{A}$. Recall that an explicit isomorphism of $\mathfrak{A}$ with a matrix ring over a division algebra has been computed in Step (4.a), and a particular form of $*$ on $\mathfrak{A}$ has been computed in Step (4.b). By these two pieces of information, we can reduce the $*$-symmetric decomposition problem for $\mathfrak{A}$ to the isometry problem for a *single* classical (symmetric, skew-symmetric, Hermitian...) form. See Section 3.4.5.

4.f. *Solve the isometry problem for a single form.* To solve the isometry problem for a single classical form is a classical algorithmic problem. One approach is to transform a given form into the standard form, by first block diagonalizing it, and then bringing the diagonal blocks to basic ones. Do this for both forms, compare whether the respective standard forms are the same, and if so, recover the isometry from the changes of bases in the standardizing procedures. See Section 3.4.6.

From Step (4.f) above, we may view the whole procedure as a reduction from isometry testing of an ϵ-symmetric matrix tuple to isometry testing of classical forms. Over $\mathbb{R}$, these classical forms are exactly those ones that define the classical groups in the sense of Weyl [Wey97] (see Section 2). In particular, in principle all possible classical forms – symmetric, skew-symmetric, Hermitian, skew-Hermitian over $\mathbb{R}$, $\mathbb{C}$, and the quaternion algebra $\mathbb{H}$ – can arise, even when we deal with only a symmetric matrix tuple, and it will be interesting to implement our algorithm and examine whether every classical form type indeed arises.

There is a tricky issue if we want to output an isometry over $\mathbb{R}$ and $\mathbb{C}$ as described in Theorem 1.4 (2) and (3). Over $\mathbb{R}$ and $\mathbb{C}$, the simple summands of a semisimple algebra may be defined over different extension fields, and one needs to be careful not to mix these fields arbitrarily as that may lead to an extension field of exponential degree. To overcome this problem we need an alternative solution to the $*$-symmetric decomposition problem as described in [IQ17, Sec. 3.5], based on $*$-invariant Wedderburn-Malcev complements of the Jacobson ideal of a $*$-algebra [Taf57].

An algorithm for Theorem 1.5 under certain technical conditions. Recall that in the ϵ-symmetrization problem, we are given a matrix tuple $\mathbf{B} = (B_1, \dots, B_m) \in M(n, \mathbb{F})^m$, and need to decide whether there exist $A, D \in \mathrm{GL}(n, \mathbb{F})$ such that $\forall i \in [m]$, AB_iD is ϵ-symmetric. Here, we present an algorithm when (1) $\mathbb{F}$ is large enough, and (2) the Jacobson radical of a matrix algebra can be computed efficiently in a deterministic way. Note that (2) holds for finite fields [Rón90] and fields of characteristic 0 [Dic23]. This algorithm follows the strategy for module isomorphism problem as used in [CIK97], and relies crucially on Lemma 1.2. We will deal with the remaining cases (a) $|\mathbb{F}|$ is large enough but we do not assume the ability to compute the Jacobson radical in Section 4.1, and (b) $|\mathbb{F}|$ is small in Section 4.2. The algorithm for (a) is obtained by associating certain projective modules to right ideals, and adapting the algorithm here to work with that concept. The algorithm for (b) follows the strategy for module isomorphism problem as used in [BL08], and relies crucially on another lemma about $*$-algebra, namely Lemma 4.1.

To start, note that if $\dim(\cap_{i\in[m]} \ker(B_i)) + \dim(\langle \cup_{i\in[m]} \mathrm{im}(B_i)\rangle) \neq n$, then $\mathbf{B}$ cannot be ϵ-symmetrizable. If $\dim(\cap_{i\in[m]} \ker(B_i)) +$

$\dim(\langle \cup_{i\in[m]} \mathrm{im}(B_i)\rangle) = n$ but $\cap_{i\in[m]} \ker(B_i) \neq \mathbf{0}$ then we can reduce to the $\cap_{i\in[m]} \ker(B_i) = \mathbf{0}$ analogously as it is done in Step (1) for the isometry problem (Section 3.1). So in the following we assume $\cap_{i\in[m]} \ker(B_i) = \mathbf{0}$ and $\langle \cup_{i\in[m]} \mathrm{im}(B_i)\rangle = \mathbb{F}^n$.

Recall that, as explained at the beginning of Section 1.3, the ϵ-symmetrization problem is equivalent to ask whether there exists $E \in \mathrm{GL}(n, \mathbb{F})$ such that $E\mathbf{B} \in S^\epsilon(n, \mathbb{F})^m$. That is, whether the matrix space $L^\epsilon(\mathbf{B}) := \{Z \in M(n, \mathbb{F}) : \forall i \in [m], ZB_i = \epsilon B_i^t Z^t\}$ contains a full-rank matrix. A linear basis $Z_1, \ldots, Z_\ell$ of $L^\epsilon(\mathbf{B})$ can be computed efficiently.

The remaining part of the algorithm is an iteration during which we maintain a matrix $Z \in L^\epsilon(\mathbf{B})$. If Z has full rank we are done. Otherwise we try all basis elements Z_i and scalars λ from a sufficiently large subset $S \subseteq \mathbb{F}$, either to obtain a matrix $Z' = Z + \lambda Z_i$ which is of higher rank than Z, or, if every such Z' is of rank no more than that of Z, conclude that Z is of the highest rank. We intend to use the following well known fact. Let $B = \begin{pmatrix} B_{11} & 0 \\ 0 & 0 \end{pmatrix}$ and $A = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}$ be $r + r'$ by $r' + r''$ block matrices where B_{11} is an r' by r' matrix of rank r' and A_{22} is a nonzero r'' by r'' matrix. Then the matrix $A + \lambda B$ has rank larger than r' for some λ from a sufficiently large set of scalars. Formally (see e.g. [IKS10, Lemma 2.2]),

Lemma 1.1. *Let $A, B \in M(r, \mathbb{F})$ and let $S \subseteq \mathbb{F}$ such that $|S| > r$. If $A \ker(B) \not\subseteq \mathrm{im}(B)$ then $\mathrm{rk}(A + \lambda B) > \mathrm{rk}(B)$ for all but at most r $\lambda \in S$.*

Unfortunately, we are unable to show – and probably it is not true in general — that Lemma 1.1 becomes applicable for Z and at least on of the basis elements Z_i when we consider $L^\epsilon(\mathbf{B})$ as it is obviously given to us (i.e., a space of n by n matrices). However, there is another representation of $L^\epsilon(\mathbf{B})$ as a matrix space in which it provably does. And this is the point where $*$-algebras enter the picture.

To see the details, assume that $\mathbf{B} = E\mathbf{B}'$ where $E \in \mathrm{GL}(n, \mathbb{F})$ and $\mathbf{B}' \in S^\epsilon(n, \mathbb{F})^m$. Since $\mathbf{B}'$ is nondegenerate, we can identify $\mathrm{Adj}(\mathbf{B}') \subseteq M(n, \mathbb{F})^{op} \oplus M(n, \mathbb{F})$ as a subalgebra of $M(n, \mathbb{F})$ by projecting to the second component (Section 2). Then $L^\epsilon(\mathbf{B}')$ is the set of $*$-symmetric elements in $\mathrm{Adj}(\mathbf{B}')$. Moreover, it is not difficult to see that $L^\epsilon(\mathbf{B}) = L^\epsilon(\mathbf{B}')E^{-1}$. The following lemma ensures that the compositon of the map $Z \mapsto ZE$ with the left multiplication action of ZE on the largest semisimple factor of $\mathrm{Adj}(\mathbf{B}')$ is a suitable representation of $L^\epsilon(\mathbf{B})$, provided that we can compute it. Its proof is given in the full version [IQ17, Sec. 4.3] of the present paper.

Lemma 1.2. *Let $\mathfrak{A}$ be a semisimple $*$-algebra over a field $\mathbb{F}$, $\mathrm{char}(\mathbb{F}) \neq 2$. Let $a \in \mathfrak{A}$ be a $*$-symmetric zero-divisor. Then there exists a $*$-symmetric element $b \in \mathfrak{A}$, such that $b\mathrm{Ann}_r(a) \not\subseteq a\mathfrak{A}$, where $\mathrm{Ann}_r(\cdot)$ denotes the set of right annihilators.*

Indeed, if b is as in Lemma 1.2 in a semisimple $\mathfrak{A}$, then viewing a and b as linear maps on $\mathfrak{A}$ (by multiplication from the left), Lemma 1.1 gives that we have that for some $\lambda \in S \subseteq \mathbb{F}$, $|S| > \dim(\mathfrak{A})$, $\dim((a + \lambda b)\mathfrak{A}) > \dim(a\mathfrak{A})$. (When working with non-semisimple algebras, we also make use the simple fact that an element of an alegbra is a unit if and only if it is a unit modulo the radical.)

Thus we wish to work with $\mathrm{Adj}(\mathbf{B}')$ and the dimension of the image of the left multiplication of its symmetric elements, that is, dimension of right ideals of the form $X\mathrm{Adj}(\mathbf{B}')$, $X \in L^\epsilon(\mathbf{B}')$ – modulo the radical of $\mathrm{Adj}(\mathbf{B}')$. But as $\mathbf{B}'$ is not in our hand, $\mathrm{Adj}(\mathbf{B}')$ and $L^\epsilon(\mathbf{B}')$ are not either. In fact $\mathbf{B}'$ is not even uniquely determined by $\mathbf{B}$. These difficulties can be overcome as follows.

- For $\mathrm{Adj}(\mathbf{B}')$, though $\mathbf{B}$ is not ϵ-symmetric, we may still define the adjoint algebra of $\mathbf{B}$ as $\mathrm{Adj}(\mathbf{B}) = \{A \oplus D \in M(n, \mathbb{F})^{op} \oplus M(n, \mathbb{F}) \mid \forall i \in [m], A^t B_i = B_i D\}$. However, while $\mathrm{Adj}(\mathbf{B}')$ is naturally a $*$-algebra by $(A \oplus D)^* = D \oplus A$, $\mathrm{Adj}(\mathbf{B})$ is not. But the following relation is easy to verify: $A \oplus D \in \mathrm{Adj}(E\mathbf{B}') \Leftrightarrow E^t A E^{-t} \oplus D \in \mathrm{Adj}(\mathbf{B}')$. So the projection of $\mathrm{Adj}(\mathbf{B})$ to the second component coincides with the projection of $\mathrm{Adj}(\mathbf{B}')$ to the second component.

- To get around the lack of $L^\epsilon(\mathbf{B}')$ is trickier. We first observe that $L^\epsilon(E\mathbf{B}F) = F^t L^\epsilon(\mathbf{B})E^{-1}$. Since $\mathbf{B} = E\mathbf{B}'$, $L^\epsilon(\mathbf{B}) = L^\epsilon(\mathbf{B}')E^{-1}$ so any $Z \in L^\epsilon(\mathbf{B})$ equals XE^{-1} for some $X \in L^\epsilon(\mathbf{B}')$. Then consider $XL^\epsilon(\mathbf{B}')$: we have $XL^\epsilon(\mathbf{B}') = XE^{-1}EL^\epsilon(\mathbf{B}') = ZL^\epsilon(\mathbf{B}'E^t) = ZL^\epsilon(\epsilon\mathbf{B}'^t E^t) = ZL^\epsilon(\epsilon(E\mathbf{B}')^t) = ZL^\epsilon(\epsilon\mathbf{B}^t)$. Here we use the assumption that $\mathbf{B}' \in S^\epsilon(n, \mathbb{F})^m$.

As $L^\epsilon(\mathbf{B}') \subseteq \mathrm{Adj}(\mathbf{B}')$, $L^\epsilon(\mathbf{B}')\mathrm{Adj}(\mathbf{B}') = \mathrm{Adj}(\mathbf{B}')$. Therefore, for any $Z \in L^\epsilon(\mathbf{B})$, $ZL^\epsilon(\epsilon\mathbf{B}^t)\mathrm{Adj}(\mathbf{B}) = XL^\epsilon(\mathbf{B}')\mathrm{Adj}(\mathbf{B}') = X\mathrm{Adj}(\mathbf{B}')$ for some $X \in L^\epsilon(\mathbf{B}')$. Noting that $L^\epsilon(\mathbf{B})$, $L^\epsilon(\epsilon\mathbf{B}^t)$, and $\mathrm{Adj}(\mathbf{B})$ are what we can compute, this allows us to work with the right ideals generated by $X \in L^\epsilon(\mathbf{B}')$ without knowing the hidden $\mathbf{B}'$.

The arguments above lead to the following algorithm, assuming that $|\mathbb{F}| > n^2$ and $J(\mathfrak{A})$ can be computed efficiently over $\mathbb{F}$. Fix $S \subseteq \mathbb{F}$ of size $> n^2$, and perform the following:

1. Compute a basis of $L^\epsilon(\mathbf{B}) = \langle Z_1, \dots, Z_\ell \rangle$, and choose some $Z \in L^\epsilon(\mathbf{B})$.

2. If Z is full-rank, return Z. Otherwise, compute $R_Z = ZL^\epsilon(\epsilon\mathbf{B}^t)\mathrm{Adj}(\mathbf{B})$.

3. If there exist $i \in [\ell]$ and $\lambda \in S$ such that $\dim(R_{Z+\lambda Z_i} + J(\mathrm{Adj}(\mathbf{B}))) > \dim(R_Z + J(\mathrm{Adj}(\mathbf{B}))$, let $Z \leftarrow Z + \lambda Z_i$ and go to Step (1). Otherwise return "Not ϵ-symmetrizable".

It is clear that the algorithm uses polynomially many arithmetic operations, and over number fields the bit sizes are controlled well. The correctness follows from Lemma 1.2: since the condition $b\mathrm{Ann}_r(a) \not\subseteq a\mathfrak{A}$ is linear, any basis of $L^\epsilon(\mathbf{B})$ contains (implicitly) such a b.

Organization of the article. In Section 2, we present certain preliminaries, including those structural results of $*$-algebras that are relevant to us. In Sections 3, we give a detailed description of the algorithm for Theorems 1.4. In Section 4, we show that for the ϵ-symmetrization problem, how to handle the cases when the Jacobson radical is not known to be efficiently computable, or the field is too small, finishing the proof of Theorem 1.5.

Due to page constraint we omit some details. The first one is the technique required to output the explicit isometry over $\mathbb{R}$ and $\mathbb{C}$ as in Theorem 1.4. The second one is the proofs of Lemmas 1.2 and 4.1. The third one is a detailed comparison with [BFP15]. They could be found in Section 3.5, Section 4.3, and Appendix, in the full version of this paper [IQ17].

2 Preliminaries

Notation. For $n \in \mathbb{N}$, $[n] := \{1, \dots, n\}$. For a field $\mathbb{F}$, $\mathrm{char}(\mathbb{F})$ denotes the characteristic of $\mathbb{F}$. $\mathbf{0}$ is the zero vector. For $B \in M(n, \mathbb{F})$, $i, j \in [n]$, $S, T \subseteq [n]$, $B(i, j)$ is the (i, j)th entry of B, $B(S, T)$ is the submatrix indexed by row indices in S and column indices in T. I_n denotes the $n \times n$ identity matrix. $\langle \cdot \rangle$ denotes the linear span.

Given a quadratic field extension $\mathbb{F}/\mathbb{F}'$, for $\alpha \in \mathbb{F}$, its conjugation $\overline{\alpha}$ is the image of α under the quadratic field involution. When $\mathbb{F} = \mathbb{C}$ and $\mathbb{F}' = \mathbb{R}$ this is simply the complex conjugation. We use $\mathbb{H}$ to denote the quaternion division algebra over $\mathbb{R}$, and i, j, k be the fundamental quaternion units. For $\alpha = a+bi+cj+dk \in \mathbb{H}$, its conjugation, denoted also by $\overline{\alpha}$, is $a-bi-cj+dk$. Given $A \in M(n, \mathbb{F})$ or $M(n, \mathbb{H})$, $\overline{A}$ denotes the matrix obtained by applying conjugation to every entry of A. For $\epsilon \in \{1, -1\}$ and $A \in M(n, \mathbb{F})$ or $M(n, \mathbb{H})$, A is ϵ-Hermitian, if $\overline{A}^t = \epsilon A$.

We will also meet matrices over division rings, and therefore, for a division ring D, the notation $M(n, D)$ (for the full $n \times n$ matrix ring over D) and $\mathrm{GL}(n, D)$ (for the group of units in $M(n, D)$).

Representation of fields and field extensions. For the isometry problem, we assume the input matrices are over a field $\mathbb{E}$ such that $\mathbb{E}$ is a finite extension of its prime field $\mathbb{F}$ (so $\mathbb{F}$ is either a field of prime order or $\mathbb{Q}$). Therefore $\mathbb{E}$ is a finite-dimensional algebra over $\mathbb{F}$. If $\dim_{\mathbb{F}}(\mathbb{E}) = d$ then $\mathbb{E}$ is the extension of $\mathbb{F}$ by a single generating element α. $\mathbb{E}$ then can be represented by the minimal polynomial of α over $\mathbb{F}$, together with an isolating interval for α in the case of $\mathbb{R}$, or an isolating rectangle for α in the case of $\mathbb{C}$. When we say that we work over $\mathbb{R}$ (resp. $\mathbb{C}$), the input is given as over a number field $\mathbb{E} \subseteq \mathbb{R}$ (resp. $\mathbb{E} \subseteq \mathbb{C}$). The algorithm is then allowed to work with extension fields of $\mathbb{E}$ in $\mathbb{R}$ (resp. $\mathbb{C}$), as long as the extension degrees are polynomially bounded. On the other hand, if we say that we work with a number field, we usually assume that we do not need to work with further extensions.

For the ϵ-symmetrization problem, we work with the arithmetic model, namely the fundamental steps are basic field operations, and the complexity is determined by counting the number of such basic operations. Furthermore, over number fields we are also concerned with the bit complexity. So when we say that some procedure works over any field, we mean that the procedure uses polynomially arithmetic operations, and when over number fields, $\mathbb{R}$ or $\mathbb{C}$, the bit complexity is also polynomial.

Tuples of matrices. A matrix tuple is an element in $M(n, \mathbb{F})^m$, and an ϵ-symmetric matrix tuple is an element in $S^\epsilon(n, \mathbb{F})^m$. We will mostly use $\mathbf{B}$, $\mathbf{C}$ to denote matrix tuples. Given $\mathbf{B} = (B_1, \dots, B_m) \in M(n, \mathbb{F})^m$, define its kernel, $\ker(\mathbf{B})$, as $\cap_{i \in [m]} \ker(B_i)$, and its image, $\mathrm{im}(\mathbf{B})$, as $\langle \cup_{i \in [m]} \mathrm{im}(B_i) \rangle$. $\mathbf{B} \in M(n, \mathbb{F})^m$ is *non-degenerate*, if $\ker(\mathbf{B}) = \mathbf{0}$, and $\mathrm{im}(\mathbf{B}) = \mathbb{F}^n$. For $\mathbf{B} \in S^\epsilon(n, \mathbb{F})^m$, due to the ϵ-symmetric condition, it can be verified easily that $\mathrm{im}(\mathbf{B}) = \{v \in \mathbb{F}^n : \forall u \in \ker(\mathbf{B}), u^t v = 0\}$. So $\mathbf{B} \in S^\epsilon(n, \mathbb{F})^m$ is non-degenerate if and only if $\ker(\mathbf{B}) = \mathbf{0}$.

Given $\mathbf{B} = (B_1, \dots, B_m) \in M(n, \mathbb{F})^m$, $\mathbf{B}^t = (B_1^t, \dots, B_m^t)$. Given $\alpha \in \mathbb{F}$, $\alpha\mathbf{B} = (\alpha B_1, \dots, \alpha B_m)$. So for $\mathbf{B} \in S^\epsilon(n, \mathbb{F})$, $\mathbf{B}^t = \epsilon\mathbf{B}$. Given $A, D \in M(n, \mathbb{F})$, $A\mathbf{B}D = (AB_1D, \dots, AB_mD)$. Given $\mathbf{B}, \mathbf{C} \in M(n, \mathbb{F})^m$, $\mathbf{B}$ and $\mathbf{C}$ are *conjugate*, if there exists $A \in \mathrm{GL}(n, \mathbb{F})$ such that $A\mathbf{B} = \mathbf{C}A$. $\mathbf{B}$ and $\mathbf{C}$ are *equivalent*, if there exists $A, D \in \mathrm{GL}(n, \mathbb{F})$ such that $A\mathbf{B} = \mathbf{C}D$. The classical module isomorphism problem asks to decide whether $\mathbf{B}$ and $\mathbf{C}$ are conjugate.

THEOREM 2.1. ([CIK97, BL08, IKS10]) *Let $\mathbf{B}$ and $\mathbf{C}$ be from $M(n, \mathbb{F})^m$. There exists a deterministic algorithm that decide whether $\mathbf{B}$ and $\mathbf{C}$ are conjugate. The algorithm uses polynomially many arithmetic operations. Over number fields the bit complexity of the* *algorithm is also polynomial.*

Structure of $*$-algebras. We collect basic facts about $*$-algebras here. A classical reference for $*$-algebras is Albert's book [Alb39]. Fix a field $\mathbb{F}$, and let $\mathfrak{A}$ be an $\mathbb{F}$-algebra, e.g. an algebra over $\mathbb{F}$. Given an anti-automorphism $* : \mathfrak{A} \to \mathfrak{A}$ of order at most 2, $(\mathfrak{A}, *)$ is termed as a $*$-algebra. We will always assume that for an $\mathbb{F}$-algebra $\mathfrak{A}$, $*$ fixes $\mathbb{F}$, that is $\alpha^* = \alpha$ for $\alpha \in \mathbb{F}$. An element $a \in \mathfrak{A}$ is *$*$-symmetric* if $a^* = a$, and *$*$-unitary* if $a^*a = 1$. A $*$-homomorphism between $(\mathfrak{A}, *)$ and $(\mathfrak{A}', \circ)$ is an algebra homomorphism $\phi : \mathfrak{A} \to \mathfrak{A}'$ such that $\phi(a^*) = \phi(a)^\circ$. An ideal $I \subseteq \mathfrak{A}$ is an $*$-ideal, if $I^* = I$. The Jacobson radical of $\mathfrak{A}$, denoted as $J(\mathfrak{A})$, is a $*$-ideal. A $*$-algebra is $*$-simple, if it does not contain non-trivial $*$-ideals. Note that for a $*$-algebra $(S, *)$, if S is simple, then it must be $*$-simple. The semisimple $\mathfrak{A}/J(\mathfrak{A})$, with the induced involution (again denoted as $*$), is $*$-isomorphic to $(S_1, *) \oplus (S_2, *) \oplus \cdots \oplus (S_k, *)$, where each $(S_i, *)$ is a $*$-simple algebra.

A $*$-simple algebra $(S, *)$ over $\mathbb{F}$ falls into two categories. Either S is a simple algebra, or S is a direct sum of two anti-isomorphic simple algebras with $*$ interchanging the two summands [Alb39, Chap. X.3]. We shall refer to the latter as *exchange type*, and its structure is simple. Specifically, recall that a simple algebra over $\mathbb{F}$ is isomorphic to $M(n, D)$ where D is a division algebra over $\mathbb{F}$. Then an exchange-type $*$-simple algebra $(S, *)$ is $*$-isomorphic to $(M(n, D) \oplus M(n, D)^{op}, \circ)$, where $\circ$ is an involution sending (A, B) to $(\phi^{-1}(B), \phi(A))$ for some algebra automorphism ϕ of $M(n, D)$.

When S is simple, a general result regarding the possible forms of involutions is [Alb39, Chap. X.4, Theorem 11]. We can explicitly list these forms for $\mathbb{F}_q$ with q odd, $\mathbb{R}$, and $\mathbb{C}$ as follows.

Over $\mathbb{F}_q$ with q odd, finite simple $*$-algebras are classified as follows (see also [BW12, Sec. 3.3]). To start with, recall that a finite simple algebra S over $\mathbb{F}_q$ is isomorphic to $M(n, \mathbb{F}_{q'})$ where $\mathbb{F}_{q'}$ is an extension field of $\mathbb{F}_q$. So without loss of generality we may assume $S = M(n, \mathbb{F}_{q'})$. Then any involution $*$ on $M(n, \mathbb{F}_{q'})$ is in one of the following forms.

- *Orthogonal type* For $X \in M(n, \mathbb{F}_{q'})$, $X^* = A^{-1}X^tA$ for some $A \in \mathrm{GL}(n, \mathbb{F}_{q'})$, $A = A^t$.
- *Symplectic type* For $X \in M(n, \mathbb{F}_{q'})$, $X^* = A^{-1}X^tA$ for some $A \in \mathrm{GL}(n, \mathbb{F}_{q'})$, $A = -A^t$.
- *Hermitian type* $\mathbb{F}_{q'}$ is a quadratic extension of a subfield $\mathbb{F}_{q''}$. For $X \in M(n, \mathbb{F}_{q'})$, $X^* = A^{-1}\overline{X}^tA$ for some $A \in \mathrm{GL}(n, \mathbb{F}_{q'})$, $\overline{A}^t = A$.

Over $\mathbb{R}$, finite simple $*$-algebras are classified as follows (see also [Lew77, Sec. 3]). To start with, recall that a finite simple algebra S over $\mathbb{R}$ is isomorphic to either $M(n, \mathbb{R})$, $M(n, \mathbb{C})$, or $M(n, \mathbb{H})$. So without loss of generality we may assume S is one of the above. Then any involution $*$ on S is in one of the following forms. Note that each type corresponds to a classical group as in [Wey97].

- *Orthogonal type* $S = M(n, \mathbb{R})$. For $X \in M(n, \mathbb{R})$, $X^* = A^{-1}X^tA$, $A \in \mathrm{GL}(n, \mathbb{R})$, $A = A^t$.
- *Symplectic type* $S = M(n, \mathbb{R})$. For $X \in M(n, \mathbb{R})$, $X^* = A^{-1}X^tA$, $A \in \mathrm{GL}(n, \mathbb{R})$, $A = -A^t$.
- *Complex orthogonal type* $S = M(n, \mathbb{C})$. For $X \in M(n, \mathbb{C})$, $X^* = A^{-1}X^tA$, $A \in \mathrm{GL}(n, \mathbb{C})$, $A = A^t$.
- *Complex symplectic type* $S = M(n, \mathbb{C})$. For $X \in M(n, \mathbb{C})$, $X^* = A^{-1}X^tA$, $A \in \mathrm{GL}(n, \mathbb{C})$, $A = -A^t$.
- *Unitary type* $S = M(n, \mathbb{C})$. For $X \in M(n, \mathbb{C})$, $X^* = A^{-1}\overline{X}^tA$, $A \in \mathrm{GL}(n, \mathbb{C})$, $A = \overline{A}^t$.
- *Quaternion unitary type* $S = M(n, \mathbb{H})$. For $X \in M(n, \mathbb{H})$, $X^* = A^{-1}\overline{X}^tA$, $A \in \mathrm{GL}(n, \mathbb{H})$, $A = \overline{A}^t$.
- *Quaternion orthogonal type* $S = M(n, \mathbb{H})$. For $X \in M(n, \mathbb{H})$, $X^* = A^{-1}\overline{X}^tA$, $A \in \mathrm{GL}(n, \mathbb{H})$, $A = -\overline{A}^t$.

On $\mathbb{C}$, $\overline{\cdot}$ denotes the standard conjugation $a + bi \mapsto a - bi$, while on $\mathbb{H}$ it is $a+bi+cj+dk \mapsto a-bi-cj-dk$.

Over $\mathbb{C}$, finite simple $*$-algebras are classified as follows. To start with, recall that a finite simple algebra S over $\mathbb{C}$ is isomorphic to $M(n, \mathbb{C})$. So without loss of generality we may assume S is $M(n, \mathbb{C})$. Then any involution $*$ on S is in one of the following forms.

- *Orthogonal type* For $X \in M(n, \mathbb{C})$, $X^* = A^{-1}X^tA$, $A \in \mathrm{GL}(n, \mathbb{C})$, $A = A^t$.
- *Symplectic type* For $X \in M(n, \mathbb{C})$, $X^* = A^{-1}X^tA$, $A \in \mathrm{GL}(n, \mathbb{C})$, $A = -A^t$.

Adjoint algebras of ϵ-symmetric matrix tuples. We first present the formal definition.

Definition 2.1. *Let $\mathbb{F}$ be a field and fix $\epsilon \in \{1, -1\}$. For $\mathbf{B} = (B_1, \ldots, B_m) \in S^\epsilon(n, \mathbb{F})^m$, the* adjoint algebra *of $\mathbf{B}$, denoted as* Adj($\mathbf{B}$), *is* $\{(A, D) \in M(n, \mathbb{F})^{op} \oplus M(n, \mathbb{F}) | \forall i \in [m], A^tB_i = B_iD\}$. Adj($\mathbf{B}$) *is a $*$-algebra over $\mathbb{F}$ with $(A, D)^* = (D, A)$.*

Note that it is a subalgebra of $M(n, \mathbb{F})^{op} \oplus M(n, \mathbb{F})$, $\mathbb{F}$ embeds in as $(\alpha I_n, \alpha I_n)$ for $\alpha \in \mathbb{F}$, and $*$ fixes $\mathbb{F}$. If $\mathbf{B}$ is non-degenerate then the projection of Adj($\mathbf{B}$) to either $M(n, \mathbb{F})^{op}$ or $M(n, \mathbb{F})$ is faithful. Therefore, *in the non-degenerate case*, we can identify $(\mathrm{Adj}(\mathbf{B}), *)$ as a subalgebra of $M(n, \mathbb{F})$ consisting of $\{D \in M(n, \mathbb{F}) \mid$

$\exists A \in M(n, \mathbb{F})$ s.t. $\forall i \in [m], A^t B_i = B_i D\}$, and for $D \in \mathrm{Adj}(\mathbf{B})$, D^* is just the (unique) solution of $\forall i \in [m], A^t B_i = B_i D$. In particular we have $A^t \mathbf{B} = \mathbf{B} A^*$.

Note that a linear basis of the adjoint algebra of a tuple of ϵ-symmetric matrices can be computed efficiently by solving a system of linear forms. The $*$-map is also easily implemented.

3 Proof of Theorem 1.4

An outline of the algorithm has been given in Section 1.5. In the following subsections, from Section 3.1 to 3.4, we give the detailed procedure, which solves completely the case of $\mathbb{F}_q$, as well as the decision version of the isometry problem for $\mathbb{R}$ and $\mathbb{C}$. The main algorithm fails to construct an explicit isometry as described in Theorem 1.4 (2) and (3). We remedy this by providing an alternative algorithm in [IQ17, Sec. 3.5], which replaces some steps of the main algorithm.

3.1 Main algorithm I: reduce to the nondegenerate case. This step works over any field. The procedure is standard but we give details here for completeness.

Recall that $\mathbf{B} \in S^\epsilon(n, \mathbb{F})^m$, as an ϵ-symmetric matrix tuple, is non-degenerate if $\ker(\mathbf{B}) = \mathbf{0}$ (Section 2). Now suppose we are given $\mathbf{B} \in S^\epsilon(n, \mathbb{F})^m$, and let $d = \dim(\ker(\mathbf{B}))$. Form a change of basis matrix $S = [v_1, \dots, v_n]$, $v_i \in \mathbb{F}^n$, such that $\{v_{n-d+1}, \dots, v_n\}$ is a basis of $\ker(\mathbf{B})$, and $\langle v_1, \dots, v_{n-d}\rangle$ is a complement subspace of $\ker(\mathbf{B})$. Then for every $i \in [m]$, $S^t B_i S = \begin{bmatrix} B_i' & 0 \\ 0 & 0 \end{bmatrix}$ where $B_i' \in S^\epsilon(n-d, \mathbb{F})$. We call $\mathbf{B}' = (B_1', \dots, B_m')$ a non-degenerate tuple extracted from $\mathbf{B}$. It is easy to show the following.

Proposition 3.1. *Given $\mathbf{B}, \mathbf{C} \in S^\epsilon(n, \mathbb{F})^m$, let $\mathbf{B}' \in S^\epsilon(\ell_1, \mathbb{F})^m$ (resp. $\mathbf{C}' \in S^\epsilon(\ell_2, \mathbb{F})^m$) be a non-degenerate tuple extracted from $\mathbf{B}$ (resp. $\mathbf{C}$). Then $\mathbf{B} \sim \mathbf{C}$ if and only if $\ell_1 = \ell_2$, and $\mathbf{B}' \sim \mathbf{C}'$.*

Since extracting a non-degenerate tuple from $\mathbf{B}$ involves only standard linear algebraic computations, this step can be performed in deterministic polynomial time. So in the following we can assume that $\mathbf{B}$ and $\mathbf{C}$ are both non-degenerate.

3.2 Main algorithm II: solve the twisted equivalence problem. This step works over any field. $\mathbf{B}, \mathbf{C} \in M(n, \mathbb{F})^m$ are twisted equivalent, if there exist $A, D \in \mathrm{GL}(n, \mathbb{F})$ such that $A^t \mathbf{B} = \mathbf{C} D$. This differs from the usual equivalence as in Definition 1.2 due to the transpose of A. But any solution (A, D) to the equivalence problem clearly gives a solution to the twisted equivalence problem by (A^t, D). The reason to introduce the twisted equivalence is because we want to be closer to the isometry concept. We now show how to test whether $\mathbf{B}$ and $\mathbf{C}$ are equivalent, by a reduction to the module isomorphism problem.

Proposition 3.2. *Given $\mathbf{B}, \mathbf{C} \in M(n, \mathbb{F})^m$, there exists a deterministic algorithm that decides whether $\mathbf{B}$ and $\mathbf{C}$ are equivalent (and therefore twisted equivalent). The algorithm uses polynomially many arithmetic operations. Over number fields the bit complexity of the algorithm is also polynomial.*

Proof. From $\mathbf{B} = (B_1, \dots, B_m)$, construct a tuple of matrices $\mathbf{B}' = (B_0', B_1', \dots, B_m')$, where $B_i' \in M(2n, \mathbb{F})$, as follows. Every B_i is viewed as a 2×2 block matrix with each block of size $n \times n$. $B_0' = \begin{bmatrix} I_n & 0 \\ 0 & 0 \end{bmatrix}$, and for $i \in [m]$, $B_i' = \begin{bmatrix} 0 & B_i \\ 0 & 0 \end{bmatrix}$. Similarly construct $\mathbf{C}'$.

We claim that there exist $A, D \in \mathrm{GL}(n, \mathbb{F})$ satisfying $A\mathbf{B} = \mathbf{C}D$ if and only if there exists an invertible $E \in \mathrm{GL}(2n, \mathbb{F})$ satisfying $E\mathbf{B}' = \mathbf{C}'E$. For the if direction, let $E = \begin{bmatrix} A & G \\ H & D \end{bmatrix}$. By $EB_0' = C_0'E$, we have $G = H = 0$. Therefore, as $E \in \mathrm{GL}(2n, \mathbb{F})$, $A, D \in \mathrm{GL}(n, \mathbb{F})$. Furthermore, for $i \in [m]$, by $\begin{bmatrix} A & 0 \\ 0 & D \end{bmatrix} \begin{bmatrix} 0 & B_i \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & C_i \\ 0 & 0 \end{bmatrix} \begin{bmatrix} A & 0 \\ 0 & D \end{bmatrix}$, we see that $AB_i = C_i D$. For the only if direction, if $AB_i = C_i D$ for all $i \in [m]$, then it is easy to see that $E = \begin{bmatrix} A & 0 \\ 0 & D \end{bmatrix}$ satisfies that $E\mathbf{B}' = \mathbf{C}'E$.

Therefore, the above construction gives an efficient reduction from the equivalence problem for $\mathbf{B}$ and $\mathbf{C}$ to the conjugacy problem for $\mathbf{B}'$ and $\mathbf{C}'$. We can then call the procedure in Theorem 2.1 to conclude.

Note that if $\mathbf{B} \sim \mathbf{C}$ then $\mathbf{B}$ and $\mathbf{C}$ are indeed twisted equivalent. In other words, if $\mathbf{B}$ and $\mathbf{C}$ are not twisted equivalent we conclude that they are not isometric either. Therefore, in the following we assume that we have computed $A, D \in \mathrm{GL}(n, \mathbb{F})$ such that $A^t \mathbf{B} = \mathbf{C} D$.

3.3 Main algorithm III: reduce to decomposing a $*$-symmetric element in a $*$-algebra. This step works over any field. From previous steps, for the non-degenerate $\mathbf{B}, \mathbf{C} \in S^\epsilon(n, \mathbb{F})$, we have computed $A, D \in \mathrm{GL}(n, \mathbb{F})$ such that $A^t \mathbf{B} = \mathbf{C} D$.

Let $\mathfrak{A} = \mathrm{Adj}(\mathbf{C})$, with the natural involution $*$. Since $\mathbf{C}$ is non-degenerate, $\mathfrak{A}$ can be embedded as a subalgebra of $M(n, \mathbb{F})$ (see Section 2.) Let $E = A^{-1} D^{-1}$. Note that E is invertible.

Claim 3.1. *Let E and $\mathfrak{A}$ be as above. E is a $*$-symmetric element in $\mathfrak{A}$.*

Proof. Observe that $A^t\mathbf{B} = \mathbf{C}D \Leftrightarrow \mathbf{B}D^{-1} = A^{-t}\mathbf{C} \Leftrightarrow D^{-t}\mathbf{B}^t = \mathbf{C}^tA^{-1} \Leftrightarrow D^{-t}\mathbf{B} = \mathbf{C}A^{-1}$, where the last $\Leftrightarrow$ uses that $\mathbf{B}$ and $\mathbf{C}$ are from $S^\epsilon(n,\mathbb{F})$. Therefore $(A^{-1}D^{-1})^t\mathbf{C} = D^{-t}A^{-t}\mathbf{C} = D^{-t}\mathbf{B}D^{-1} = \mathbf{C}A^{-1}D^{-1}$.

The following proposition is a conceptually crucial observation for the algorithm.

Proposition 3.3. *Let $\mathbf{B}$, $\mathbf{C}$, $\mathfrak{A}$, and E be as above. Then $\mathbf{B}\sim\mathbf{C}$ if and only if there exists $X\in\mathfrak{A}$ such that $X^*X = E$.*

Proof. For the if direction, by $X^*X = A^{-1}D^{-1}$, we have $AX^* = D^{-1}X^{-1}$. Also observe that $D^{-t}\mathbf{B} = \mathbf{C}A^{-1}$, and $(X^*)^t\mathbf{C} = \mathbf{C}X \Leftrightarrow \mathbf{C}X^* = X^t\mathbf{C} \Leftrightarrow X^{-t}\mathbf{C} = \mathbf{C}(X^*)^{-1}$. So $(D^{-1}X^{-1})^t\mathbf{B} = X^{-t}D^{-t}\mathbf{B} = X^{-t}\mathbf{C}A^{-1} = \mathbf{C}(X^*)^{-1}A^{-1}$, which gives $(D^{-1}X^{-1})^t\mathbf{B}(AX^*) = \mathbf{C}$. Now recall that $AX^* = D^{-1}X^{-1}$, so $D^{-1}X^{-1}$ is the desired isometry.

For the only if direction, suppose $Z^t\mathbf{B}Z = \mathbf{C}$. Setting $X = Z^{-1}D^{-1}$ and $Y = A^{-1}Z$, we have $AY = D^{-1}X^{-1} = Z$. So $\mathbf{C} = Z^t\mathbf{B}Z = Y^tA^t\mathbf{B}D^{-1}X^{-1} = Y^t\mathbf{C}X^{-1}$, which gives $Y = X^*$. By $YX = A^{-1}D^{-1}$, $X^*X = A^{-1}D^{-1}$ follows.

Proposition 3.3 then leads to the following question.

Problem 3.1. ($*$-symmetric decomposition problem) *Let $\mathfrak{A}$ be a matrix algebra in $M(n,\mathbb{F})$ with an involution $*$, and $E\in\mathfrak{A}$ be an invertible $*$-symmetric element. Compute $X\in\mathfrak{A}$ such that $X^*X = E$, if there exists such an element.*

3.4 Main algorithm IV: solve the $*$-symmetric decomposition problem. This is the main technical piece of this algorithm. The strategy is to utilize the algebra structure of $\mathfrak{A}$, and reduce the problem to the case when $\mathfrak{A}$ is a simple algebra. When $\mathfrak{A}$ is simple and can be explicitly represented as a full matrix ring over division algebras, the problem turns out to be equivalent to solving the isometry problem for a single classical (symmetric, skew-symmetric, Hermitian...) form, which then can be solved using existing algorithms.

3.4.1 Decomposition algorithm I: compute the algebra structure. By resorting to known results, this step works over finite fields [Rón90, Iva00, EG00], the real field, and the complex field [FR85,Ebe91a,Ebe91b]. We now cite these results as follows.

Theorem 3.1. ([Rón90]; see also [Iva00, EG00]) *Suppose we are given a linear basis of an algebra $\mathfrak{A}$ in $M(n,\mathbb{F}_q)$. There is a Las Vegas algorithm that computes*

1. *a linear basis of the Jacobson radical $J(\mathfrak{A})$, and*
2. *an epimorphism $\pi:\mathfrak{A}\to M(n_1,\mathbb{F}_{q_1})\oplus\cdots\oplus M(n_k,\mathbb{F}_{q_k})$ with kernel $J(\mathfrak{A})$, and $\mathbb{F}_{q_i}$ an extension field of $\mathbb{F}_q$. $\mathbb{F}_{q_i}$ is specified by a linear basis over $\mathbb{F}_q$.*

The algorithm runs in time $\mathrm{poly}(n,\log q)$, *and can be derandomized at the price of running in time* $\mathrm{poly}(n,\log q,p)$ *where* $p = \mathrm{char}(\mathbb{F}_q)$.

Furthermore, there are efficient deterministic algorithms that

i. *given $a\in\mathfrak{A}$, compute $\pi(a)$, and*

ii. *given $b\in M(n_1,\mathbb{F}_{q_1})\oplus\cdots\oplus M(n_k,\mathbb{F}_{q_k})$, compute $a\in\mathfrak{A}$ such that $\pi(a)=b$.*

Theorem 3.2. ([FR85, Ebe91a, Ebe91b, Rón94]) *Let $\mathbb{E}$ be a number field, and suppose we are given a linear basis of an algebra $\mathfrak{A}$ in $M(n,\mathbb{E})$. Then there exists a deterministic polynomial-time algorithm that computes*

1. *a linear basis of the Jacobson radical $J(\mathfrak{A})$ over $\mathbb{E}$, and*
2.
   - *Over $\mathbb{R}$: (a) the number k of simple components of $\mathfrak{A}\otimes_{\mathbb{E}}\mathbb{R}$,*

     (b) specifications of extension fields $\mathbb{E}\subseteq\mathbb{E}_1,\dots,\mathbb{E}_k\subseteq\mathbb{R}$, such that each $\mathbb{E}_i$ is of degree at most $\binom{\dim_{\mathbb{E}}\mathfrak{A}}{2}$ over $\mathbb{E}$,

     (c) bases of simple algebras $B_1\subseteq A\otimes_{\mathbb{E}}\mathbb{E}_1$, ..., $B_k\subseteq A\otimes_{\mathbb{E}}\mathbb{E}_k$, such that $B_i\otimes_{\mathbb{E}_i}\mathbb{R}$, $i\in[k]$, are all the simple components of $\mathfrak{A}\otimes_{\mathbb{E}}\mathbb{R}$, and

     (d) for each $i\in[k]$, an extension field $\mathbb{K}_i\subseteq\mathbb{R}$ over $\mathbb{E}_i$ with extension degree at most $\dim_{\mathbb{E}_i}B_i$, the linear basis of a division algebra $D_i\subseteq B_i\otimes_{\mathbb{E}_i}\mathbb{K}_i$ over $\mathbb{K}_i$, and the linear basis of a subalgebra $M_i\subseteq B_i\otimes_{\mathbb{E}_i}\mathbb{K}_i$ over $\mathbb{K}_i$, such that $M_i\cong M(n_i,\mathbb{K}_i)$, and $B_i\otimes_{\mathbb{E}_i}\mathbb{K}_i\cong M_i\otimes_{\mathbb{K}_i}D_i\cong M(n_i,D_i)$. $\dim_{\mathbb{K}_i}D_i$ can be 1, 2, or 4, and when $\dim_{\mathbb{K}_i}D_i=4$, D_i is non-commutative.
   - *Over $\mathbb{C}$: (a) the number k of simple components of $\mathfrak{A}\otimes_{\mathbb{E}}\mathbb{C}$,*

     (b) specifications of extension fields $\mathbb{E}\subseteq\mathbb{E}_1,\dots,\mathbb{E}_k$, such that each $\mathbb{E}_i$ is of degree at most $\dim_{\mathbb{E}}\mathfrak{A}$ over $\mathbb{E}$,

     (c) bases of simple algebras $B_1\subseteq A\otimes_{\mathbb{E}}\mathbb{E}_1$, ..., $B_k\subseteq A\otimes_{\mathbb{E}}\mathbb{E}_k$, such that $B_i\otimes_{\mathbb{E}_i}\mathbb{C}$, $i\in[k]$, are all the simple components of $\mathfrak{A}\otimes_{\mathbb{E}}\mathbb{C}$, and

(d) *for each $i \in [k]$, an extension field $\mathbb{K}_i$ over $\mathbb{E}_i$ with extension degree at most $\sqrt{\dim_{\mathbb{E}_i} B_i}$, the linear basis of a subalgebra $M_i \subseteq B_i \otimes_{\mathbb{E}_i} \mathbb{K}_i$ over $\mathbb{K}_i$, such that $M_i \cong M(n_i, \mathbb{K}_i)$.*

REMARK 3.1. *1. Comparing Theorem 3.1 and Theorem 3.2, we see that a statement corresponding to Theorem 3.2 (ii) was missing in Theorem 3.1. This is because a preimage of $b \in M(n_1, D_1) \oplus \cdots \oplus M(n_k, D_k)$ may live in $\mathfrak{A} \otimes_{\mathbb{E}} \mathbb{K}$ for some field $\mathbb{K}$ with an exponential extension degree over $\mathbb{E}$. This suggests that representing the isometry in the settings of $\mathbb{R}$ and $\mathbb{C}$ as a* single *matrix would be inefficient.*

2. The randomized version of Theorem 3.2 is shown by Eberly in [Ebe91a, Ebe91b], and is subsequently derandomized by Rónyai in [Rón94]. To completely derandomize Theorem 3.1 is a difficult problem as this relies on algorithms for polynomial factorization over finite fields.

3.4.2 Decomposition algorithm II: recognize the $*$-algebra structure. This step works over $\mathbb{F}_q$ with q odd, $\mathbb{R}$, and $\mathbb{C}$. It may be possible to handle fields of even characteristics, but we leave it for further study. The case of finite fields of odd characteristics has been settled by Brooksbank and Wilson in [BW12]. Here we provide a unified and somewhat simpler treatment over those fields just mentioned.

To start with, recall that from previous steps we have computed the algebra structure of $\mathfrak{A} \subseteq M(n, \mathbb{F})$, including a linear basis of $J(\mathfrak{A})$ and an epimorphism $\pi : \mathfrak{A} \to S_1 \oplus \cdots \oplus S_k$ where S_i is a simple algebra over the designated field (after some scalar extension when over $\mathbb{R}$ or $\mathbb{C}$). We have also computed explicit isomorphisms between S_i and matrix rings over division rings. Since $J(\mathfrak{A})$ is a $*$-ideal, the involution $*$ induces an involution, which we denote again by $*$, on $\pi(\mathfrak{A})$. Then for each S_i, either $S_i^* = S_i$, or $S_i^* = S_j$ for some $j \neq i$. The goal is that, in the former case, we want to express the involution $*$ explicitly in the forms presented in Section 2.

PROPOSITION 3.4. *Let $\mathbb{E}/\mathbb{F}$ be a field extension specified by a linear basis over $\mathbb{F}$. Given an involution $*$ of $M(n, \mathbb{E})$ as an $\mathbb{F}$-algebra, there exists a deterministic polynomial-time algorithm that (1) decides whether $*$ induces a quadratic field involution of $\mathbb{E}$ over a subfield $\mathbb{E}'$, and (2) computes $A \in \mathrm{GL}(n, \mathbb{E})$ such that for every $X \in M(n, \mathbb{E})$, $X^* = A^{-1} X'^t A$, where X' is either X (when $*$ fixes $\mathbb{E}$) or $\overline{X}$ (when $*$ induces a quadratic field involution).*

Proof. For (1), we apply $*$ to every basis element b in the linear basis of $\mathbb{E}$ over $\mathbb{F}$. If $*$ changes none of them, then $\mathbb{E}$ is also invariant under $*$. If $*$ changes some of them, the sums $b + b^*$ linearly span a subfield $\mathbb{E}'$ such that $\mathbb{E}/\mathbb{E}'$ is a quadratic field extension, and $*$ induces the quadratic field involution. For (2), for any $X \in M(n, \mathbb{E})$ let X' be as defined in the statement. We take a linear basis $\{B_1, \ldots, B_{n^2}\}$ of $M(n, \mathbb{E})$ (the standard basis will do), and set up $Y B_i^* = B_i'^t Y$, for $i \in [n^2]$, and Y is an $n \times n$ variable matrix. By [Alb39, Chap. X.4, Theorem 11], there must exist some $A \in \mathrm{GL}(n, \mathbb{E})$ as a valid solution to Y in the above equations. From the algorithmic viewpoint, this is an instance of the module isomorphism problem, and we can apply the procedure in Theorem 2.1 to conclude.

Note that Proposition 3.4 covers all simple types over $\mathbb{F}_q$ with q odd and $\mathbb{C}$, as well as those simple types over $\mathbb{R}$ except the two quaternion types. We now handle the two quaternion types in the real field setting.

PROPOSITION 3.5. *Let $\mathbb{H}$ be given by a linear basis over $\mathbb{R}$. Given an involution $*$ of $M(n, \mathbb{H})$ as an $\mathbb{R}$-algebra, there exists a deterministic polynomial-time algorithm that computes $A \in \mathrm{GL}(n, \mathbb{H})$ such that for every $X \in M(n, \mathbb{H})$, $X^* = A^{-1} \overline{X}^t A$.*

Proof. Let $f : \mathbb{H} \to M(4, \mathbb{R})$ be the regular representation of $\mathbb{H}$ on $\mathbb{R}^4$. Let $\{C_1', C_2', C_3', C_4'\}$ be a linear basis of the centralizing algebra of $f(\mathbb{H})$ in $M(4, \mathbb{R})$, which is isomorphic to $\mathbb{H}^{op}$. Now think of matrices in $M(4n, \mathbb{R})$ as $n \times n$ block matrices with each block of size 4×4. For $i \in [4]$, let $C_i \in M(4n, \mathbb{R})$ be the diagonal block matrix, with all diagonal blocks being C_i'. f naturally embeds $M(n, \mathbb{H})$ to $M(4n, \mathbb{R})$. By the double centralizer theorem, the centralizing algebra of C_i's is $f(M(n, \mathbb{H}))$.

The above reasoning suggests the following construction. Take a basis $\{B_1, \ldots, B_{n^2}\}$ of $M(n, \mathbb{H})$, and let $B_i' = \overline{B_i}^t$. Set up $Y f(B_i) = f(B_i') Y$, $i \in [n^2]$, $Y C_j = C_j Y$, $j \in [4]$, where Y is a $4n \times 4n$ variable matrix. By $Y C_j = C_j Y$, any valid solution to Y lies in $f(M(n, \mathbb{H}))$. By an analogous argument as in the proof of Proposition 3.4, there must exist an invertible $A \in \mathrm{GL}(4n, \mathbb{R})$ as a valid solution to Y, and can be solved as as an instance of the module isomorphism problem by Theorem 2.1. Finally, after getting such an $A \in \mathrm{GL}(4n, \mathbb{R})$, it is straightforward to compute the preimage of A in $M(n, \mathbb{H})$, concluding the proof.

3.4.3 Decomposition algorithm III: reduce to the semisimple case. This step works over fields of characteristic $\neq 2$, and is the main bottleneck for handling fields of characteristic 2.

Proposition 3.6. *Let $\mathfrak{A}$ be a $*$-algebra over $\mathbb{F}$, $\mathrm{char}(\mathbb{F}) \neq 2$. Let $E \in \mathfrak{A}$ be an invertible $*$-symmetric element, and suppose there exists $X \in \mathfrak{A}/J(\mathfrak{A})$, such that $Y^*Y + J(\mathfrak{A}) = E + J(\mathfrak{A})$. Then there exists $X \in \mathfrak{A}$ such that $X^*X = E$, and there exists a deterministic polynomial-time algorithm that outputs such an X.*

Proof. To recover $X \in \mathfrak{A}$ such that $X^*X = E$, consider the following situation: suppose we have a $*$-ideal J of $\mathfrak{A}$ with $J^2 = 0$, and an invertible $E \in \mathfrak{A}$ with $E^* = E$. Given Y such that $Y^*Y + J = E + J$, the goal is to find $Z \in J$ such that $(Y+Z)^*(Y+Z) = E$. Expanding to $Y^*Y + Y^*Z + YZ^* + Z^*Z = E$, by $Z^*Z = 0$ we need to satisfy $Y^*Z + Z^*Y = E - Y^*Y$. Note that $E - Y^*Y$ is $*$-symmetric. So setting $U = \frac{1}{2}(E - Y^*Y)$, $Z = Y^{-*}U$ is the desired, and $X = Y + Z$ satisfies $X^*X = E$. Using this procedure we can upgrade a solution mod $J(\mathfrak{A})$ to a solution mod $J(\mathfrak{A})^2$, $J(\mathfrak{A})^4$, etc., to upgrade a solution for $Y^*Y + J(\mathfrak{A}) = E + J(\mathfrak{A})$ to a solution for $X^*X = E$.

3.4.4 Decomposition algorithm IV: reduce to the $*$-simple and simple case.

This step works for any field. Suppose we have a semi-simple algebra $\mathfrak{A}$ decomposed into a direct sum of simple summands $S_1 \oplus \cdots \oplus S_k$, and let $*$ be an involution on $\mathfrak{A}$. Without loss of generality, we can assume that there exists $j \leq \lfloor k/2 \rfloor$, such that $*$ exchanges S_{2i-1} and S_{2i} for $i \in [j]$, and stabilizes S_i for $i > 2j$. Let $E \in \mathfrak{A}$ be an invertible $*$-symmetric element, and let E_i be the projection of E to S_i. Recall that our goal is to find $X \in \mathfrak{A}$ such that $X^*X = E$, if such an X exists.

Proposition 3.7. *Let $\mathfrak{A}$, S_i, E, and E_i be as above. There exists $X \in \mathfrak{A}$ such that $X^*X = E$, if and only if for every $i > 2j$, there exists $X_i \in S_i$ such that $X_i^*X_i = E_i$.*

Proof. For the if direction, we claim that $X = E_1 \oplus I \oplus E_3 \oplus I \oplus \cdots \oplus E_{2j-1} \oplus I \oplus X_{2j+1} \oplus \cdots \oplus X_k$ is a solution, where I denotes the identity element in the respective summand. To see this, let us suppose $*$ exchanges S_1 and S_2. Then by $(E_1, E_2)^* = (E_1, E_2)$, we have $(E_1, I)^* = (I, E_2)$. So $X^*X = (I \oplus E_2 \oplus I \oplus E_4 \oplus \cdots \oplus I \oplus E_{2j} \oplus X_{2j+1}^* \oplus \cdots \oplus X_k^*)(E_1 \oplus I \oplus E_3 \oplus I \oplus \cdots \oplus E_{2j-1} \oplus I \oplus X_{2j+1} \oplus \cdots \oplus X_k) = E$.

For the only if direction, suppose $X = X_1 \oplus X_2 \oplus \cdots \oplus X_{2j-1} \oplus X_{2j} \oplus X_{2j+1} \oplus \cdots \oplus X_k$ satisfies $X^*X = E$. Then it is straightforward to verify that for $i > 2j$, $X_i^*X_i = E_i$.

3.4.5 Decomposition algorithm V: the simple case by reducing to the isometry problem for a single form.

This step works over any field. From previous steps, we now have (1) $M(n, D)$ where D is a field or a division algebra, (2) an involution $*$ on $M(n, D)$, which induces an involution $\bar{\cdot} : D \to D$ (possibly identity), such that $X^* = A^{-1}\overline{X}^t A$ and $\overline{A}^t = \epsilon A$ for some $\epsilon \in \{1, -1\}$, and (3) an invertible $*$-symmetric element E.

Here is the other conceptually crucial observation.

Proposition 3.8. *Let notation be as above. Let $F = AE$. Then F is a form of the same type as A, and there exists X such that $X^*X = E$, if and only if A and F are isometric.*

Proof. To see that F is a form of the same type as A, we have $E = E^* = A^{-1}\overline{E}^t A$ (by the $*$-symmetry of E) and $\overline{A}^t = \epsilon A$. So $AE = \overline{E}^t A$, which is equivalent to, by taking conjugate transpose, $\overline{E}^t A = AE$. Therefore $\overline{AE}^t = \overline{E}^t\overline{A}^t = \epsilon\overline{E}^t A = \epsilon AE$.

For the second statement, we consider the if direction first. If for some $Y \in \mathrm{GL}(n, D)$, $Y^t A \overline{Y} = F = AE$, then $A^{-1}Y^t A\overline{Y} = E$. Setting $X = \overline{Y}$, we have $A^{-1}\overline{X}^t AX = E$. Noting that $A^{-1}\overline{X}^t A = X^*$, we obtain the desired $X^*X = E$. The only if direction can be seen easily by inverting the above reasoning.

3.4.6 Decomposition algorithm VI: solve the isometry problem for a single form.

To solve the isometry problem for a single form over a division ring, we will in fact compute the canonical form for such a form. The isometry problem can then be solved by comparing the canonical forms. Over $\mathbb{F}_q$ with q odd, a concrete isometry can be obtained by using the transformations to the canonical forms. To recover a concrete isometry (represented in some form) over $\mathbb{R}$ or $\mathbb{C}$ requires more technical machinery and we leave it to [IQ17, Sec. 3.5]. The existence of canonical forms is well-known for $\mathbb{F}_q$ with q odd (see e.g. [Wil09c, Chap. 3.4]), for $\mathbb{R}$ (see e.g. [Lew77, Sec. 4]), and for $\mathbb{C}$.

Computing the canonical form involves two steps. Let $E \in M(n, D)$ such that $\overline{E}^t = \epsilon E$, where $\bar{\cdot} : D \to D$ is an involution, and $\epsilon \in \{1, -1\}$.

The first step is to compute an orthogonal basis for E, that is a linear basis of D^n $\{e_1, \ldots, e_n\}$, such that for every $i \in [n]$, $e_i^t E \overline{e_j} = 0$ for exactly one e_j. This is known as the Gram-Schmidt procedure, and an efficient algorithm in this general setting has been obtained by Wilson.

Theorem 3.3. ([Wil13]) *Let E be as above. There exists a deterministic polynomial-time algorithm that computes an orthogonal basis for E.*

After the first step, by transforming to the orthogonal basis, E can be assumed to be a diagonal block matrix, with each block is of size 1 or 2. The second step is to simplify these diagonal blocks as much as possible. We now need to handle each field separately. Recall that E is non-degenerate.

Block diagonal forms over $\mathbb{F}_q$. We distinguish among the three simple types over $\mathbb{F}_q$.

- *Orthogonal type* In this case, each block is of size 1, e.g. E is a diagonal matrix. Fix a non-square ω in $\mathbb{F}_q$, which can be computed efficiently, by either using randomness, or in a deterministic way if we assume the Generalized Riemann Hypothesis or the characteristic of $\mathbb{F}_q$ is small. We can first simplify E as $\text{diag}(1,\ldots,1,\omega,\ldots,\omega)$, by resorting to square root computations over finite fields. This can be done in randomized polynomial time by e.g. the Tonelli-Shanks algorithm. A deterministic polynomial-time algorithm exists, if we assume the Generalized Riemann hypothesis, or the characteristic of the finite field is small. Then, if the number of ω's is larger than 1, then write ω as a sum of two squares $\alpha^2+\beta^2$, which is always possible over a finite field. Algorithmically, this can be done by solving the equation $x^2+y^2=\omega$ in deterministic polynomial time by an algorithm of van de Woestijne [vdW05, Theorem A.3]. Given such α,β, $\text{diag}(\omega,\omega)$ can be transformed to $\text{diag}(1,1)$ by $\begin{bmatrix}\alpha & \beta\\ \beta & -\alpha\end{bmatrix}\begin{bmatrix}1 & 0\\ 0 & 1\end{bmatrix}\begin{bmatrix}\alpha & \beta\\ \beta & -\alpha\end{bmatrix}^t=\begin{bmatrix}\omega & 0\\ 0 & \omega\end{bmatrix}$. Therefore the possible standard forms are $\text{diag}(1,\ldots,1)$ or $\text{diag}(1,\ldots,1,\omega)$.
- *Symplectic type* In this case, each block is of size 2, so we examine one block $\begin{bmatrix}0 & \alpha\\ -\alpha & 0\end{bmatrix}$. Now by expressing α as a sum of squares, similar trick applies to bring it to $\begin{bmatrix}0 & 1\\ -1 & 0\end{bmatrix}$.
- *Hermitian type* In this case, each block is of size 1. Let the associated field extension be $\mathbb{F}_q/\mathbb{F}_{q'}$ where $q=q'^2$, and suppose $\mathbb{F}_q=\mathbb{F}_{q'}(\omega)$. Then for $\alpha=a+b\omega$, $\overline{\alpha}=a-b\omega$. For a diagonal entry $\alpha\in\mathbb{F}_q$, $\alpha=\overline{\alpha}$, we need to compute $\beta\in\mathbb{F}_q$ such that $\beta\overline{\beta}=\alpha$, which always exists. Setting $\beta=x+y\omega$, we need to solve the equation $\beta\overline{\beta}=x^2-y^2\omega^2=\alpha$. Again this can be solved in deterministic polynomial time by [vdW05, Theorem A.3].

Block diagonal forms over $\mathbb{R}$. For the symplectic, complex orthogonal, complex symplectic, quaternion orthogonal types, we can always bring a given form to the identity matrix or the standard non-degenerate skew-symmetric matrix. For other types, we can bring a given form to $\text{diag}(1,\ldots,1,-1,\ldots,-1)$, where the number of 1's and the number of -1's is called the signature of the canonical form. Therefore, if we just want to compare whether $\mathbf{B}$ and $\mathbf{C}$ are isometric over $\mathbb{R}$, up to this point, the signatures are the only things to compare.

Block diagonal forms over $\mathbb{C}$. For the two types here we can always bring a given form to the identity matrix or the standard non-degenerate skew-symmetric matrix. In particular, this suggests that any two forms of the same type are always isometric, so testing isometry between $\mathbf{B}$ and $\mathbf{C}$ only depends on whether $\mathbf{B}$ and $\mathbf{C}$ are equivalent or not.

4 Proof of Theorem 1.5

4.1 When $|\mathbb{F}|$ is large enough. Suppose $|\mathbb{F}|=\Omega(n^3)$. We shall extend the algorithm in Section 4 to work without relying on the presence of the radical of $\text{Adj}(\mathbf{B})$. To that end we require the following proposition.

PROPOSITION 4.1. *Let $\mathfrak{A}$ be a finite dimensional algebra with identity and let J be a non-nilpotent right ideal $\mathfrak{A}$. Then in deterministic polynomial time one can compute a right ideal J_0 contained in J generated by an idempotent e such that $e+\text{Rad}(\mathfrak{A})$ is a left identity element of $(J+\text{Rad}(\mathfrak{A}))/\text{Rad}(\mathfrak{A})$. Any such J_0, as a right $\mathfrak{A}$-module is the projective cover of the semisimple right $\mathfrak{A}$-module $(J+\text{Rad}(\mathfrak{A}))/\text{Rad}(\mathfrak{A})$ and hence depends only on the structure of J modulo $\text{Rad}(\mathfrak{A})$.*

Proof. For the last statement, see [Pie82], Section 6.4. To compute J_0 it is sufficient to find an idempotent e of J with the property as in the statement. As J is not nilpotent one can find a non-nilpotent element and even an idempotent e in J. Compute the right ideal $J''=\{x-ex : x\in J\}$. Obviously $e\mathfrak{A}\cap J''=0$. If J'' is nilpotent then e is as requested. Otherwise find an idempotent f in J''. We have $ef=0$ and $(e+fe)^2=ee+fefe+efe+fee=e+fe$, whence if $fe\neq 0$ then we can replace e with $e+fe$ which generates a right ideal larger than $e\mathfrak{A}$. If $fe=0$ then $(e+f)^2=e+f$ whence we can proceed with $e+f$ in place of e. To keep sizes moderate, we cvan slighly modify the procedure. We express e in terms of a basis for J. Let $n=\dim J$. Then for every element x, the multiplicity of zero in the minimal polynomial x^n has multiplicity at most one. Using the method of [dGIR96, Lemma 2.2], we find an element x having "small" coefficients in terms of the basis of J with the property that x^n is of rank at least as large as that of e. Then replace e with the maximal idempotent e' of the subalgebra generated by x^n (This algebra is spanned by $x^n,x^{2n},\ldots,x^{n^2}$.) By the property of x^n, e' has rank at leats as large as that of x^n. We replace e with e' and increase its rank if $(1-e)J$ is not nilpotent.

We call the right ideal J_0 as in Proposition 4.1 the projective module associated to J and denote it by $P(J)$. For a nilpotent right ideal J we set $P(J) = 0$.

Fact 4.1. *Let $\mathfrak{A}$ be a finite dimensional semisimple algebra with identity, let $\mathfrak{A}_1, \ldots, \mathfrak{A}_\ell$ be the simple components of $\mathfrak{A}$, and let $\pi_j : \mathfrak{A} \to \mathfrak{A}_j$, $j \in [\ell]$, be the corresponding projections. Suppose that e and f are idempotents in $\mathfrak{A}$ such that the rank of $\pi_j(e)$ is the same as that of $\pi_j(f)$ for $j = 1, \ldots, \ell$. Then $e\mathfrak{A}$ and $f\mathfrak{A}$ are isomorphic as right $\mathfrak{A}$-modules.*

Now we are ready to upgrade the algorithm in Section 4 to work without $J(\mathrm{Rad}(\mathbf{B}))$.

Proposition 4.2. *Let $\mathfrak{A}$ be a finite dimensional algebra with identity, let a be a zero-divisor in $\mathfrak{A}$ and let $b \in \mathfrak{A}$ such that $a + \mathrm{Rad}(\mathfrak{A})$ and $b + \mathrm{Rad}(\mathfrak{A})$ behave like a and b in Lemma 1.2. If S is a sufficiently large subset of the base field, then for at least one $\lambda \in S$ we have $\dim P((a + \lambda b)\mathfrak{A}) > \dim P(a\mathfrak{A})$.*

Proof. We have that, modulo $\mathrm{Rad}(\mathfrak{A})$, $a + \lambda b$ generates a right ideal that has dimension higher than that of generated by a for at least one λ from S if S is sufficiently large ($|S| = \Omega(n^2)$). If S is even larger ($|S| = \Omega(n^3)$), then S will contain such a λ with the additional property that for the projection of $(a+\lambda b)\mathfrak{A}$ to any of the simple components of $\mathfrak{A}/\mathrm{Rad}(\mathfrak{A})$ has dimension at least as high as that for the projection of $a\mathfrak{A}$. Then, by Fact 4.1, the right $\mathfrak{A}$-module $a\mathfrak{A}+\mathrm{Rad}(\mathfrak{A})$ can be embedded into as a proper submodule. By monotonicity of taking projective covers of semisimple modules, $P(a\mathfrak{A})$ is isomorphic to a proper submodule of $P(b\mathfrak{A})$.

4.2 When $|\mathbb{F}|$ is small. The algorithm in Section 1.5, upgraded in Section 4.1, runs in polynomial time even over a number field, but has the disadvantage of relying on the field to be large enough. In this subsection, we present an algorithm that works even for small fields. However, the disadvantage of this algorithm is that, over a number field it seems difficult to bound the bit sizes of intermediate data. Still, combining these two algorithms together we are able to cover all fields, so this proves Theorem 1.5.

As explained in Section 1.5, w.l.o.g. we can assume $\mathbf{B}$ to be non-degenerate. The following Lemma 4.1 is the key to this algorithm. Its proof can be found in [IQ17, Sec. 4.3].

Lemma 4.1. *Let $\mathbb{F}$ be a field of characteristic not 2. Let $\mathfrak{A}$ be a finite dimensional $*$-algebra over $\mathbb{F}$ with an identity element. Let a be a $*$-symmetric element of $\mathfrak{A}$ such that the right ideal $a\mathfrak{A}$ has a left identity element. Then the right annihilator $\mathrm{Ann}_r(a) = \{b \in \mathfrak{A} : ab = 0\}$ of a is generated, as a right ideal, by a $*$-symmetric element of $\mathfrak{A}$.*

We shall only sketch the idea behind the algorithm in the following; a rigorous algorithm can be extracted without much difficulty.

Suppose $\mathbf{B} = E\mathbf{B}'$ where $E \in \mathrm{GL}(n, F)$ and $\mathbf{B}' \leq S^\epsilon(n, \mathbb{F})$. We claim that $L^\epsilon(\mathbf{B}')$ cannot be spanned by nilpotent elements. Indeed, assume the contrary. Let $\mathfrak{A} = \mathrm{Adj}(\mathbf{B}')$, which is a $*$-algebra as $\mathbf{B}'$ is ϵ-symmetric. Then $I \otimes *$ is an involution of $\overline{\mathfrak{A}} = \overline{\mathbb{F}} \otimes_{\mathbb{F}} \mathfrak{A}$, where $\overline{\mathbb{F}}$ is an algebraic closure of $\mathbb{F}$. We identify $\mathfrak{A}$ with the subalgebra $1 \otimes \mathfrak{A}$ and use $*$ for $I \otimes *$. The $*$-symmetric elements of $\overline{\mathfrak{A}}$ are $\overline{\mathbb{F}}$-linear combinations of $*$-symmetric elements of $\mathfrak{A}$. Using this, we may assume that $\mathbb{F}$ is algebraically closed. Then the $*$-simple components of the factor of $\mathfrak{A}/\mathrm{Rad}(\mathfrak{A})$ contain $*$-symmetric idempotents whose images are rank one or two matrices under some irreducible representation of $\mathfrak{A}$. It follows that any basis for $L^\epsilon(\mathbf{B}')$ contains an element whose image under a matrix representation of $\mathfrak{A}$ has nonzero trace. Such an element cannot be nilpotent.

Thus any basis of $L^\epsilon(\mathbf{B}) = L^\epsilon(\mathbf{B}')E^{-1}$ contains an element of the form $Z = XE^{-1}$ where X is a non-nilpotent element of $L^\epsilon(\mathbf{B}')$. Now consider the inner ideal $XL^\epsilon(\mathbf{B}')X$. This set equals the set of the $*$-symmetric elements of the subalgebra $X\mathfrak{A}X$. This subalgebra is not nilpotent. Therefore, just like above, as it contains the non-nilpotent element X, an arbitrary basis for $XL^\epsilon(\mathbf{B}')X$ contains a non-nilpotent element. It follows that an arbitrary basis for $L^\epsilon(\mathbf{B}')$ (which may differ from the basis which X is chosen from) contains an element Y such that XYX is not nilpotent. In particular, a basis for $L^\epsilon(\mathcal{B}^*) = EL^\epsilon(\mathbf{B}')$ contains an element Z' of the form $Z' = EY$ where XYX is not nilpotent. Now consider the sequences $X_k = X(YX)^k$ and $Y_k = Y(XY)^k$, $k \geq 0$. We have $X_0 = X$, $Y_0 = Y$, $X_{k+1} = XY_kX$ and $Y_{k+1} = YX_kY$. Furthermore $X_{k+1}E^{-1} = (XE^{-1})(EY_k)(XE^{-1})$ and $EY_{k+1} = (EY)(X_kE^{-1})(EY)$, which gives an efficient method for computing X_kE^{-1} and EY_k. The kernels of X_k form a nondecreasing chain of linear spaces. Therefore if k is large enough then $\ker X_\ell = \ker X_k$ for $\ell > k$. The sequences consisting of the kernels of Y_k as well as those consisting of the images of X_k and the images of Y_k stabilize as well. From $X_{2k} = X_kY_kX$ we infer that for sufficiently large k the kernel of X_kY_k is the same as that of X_k and the image of X_kY_k is the same as that of Y_k. Analogous equalities hold for the kernel and for the image of Y_kX_k. These properties of the pair X_k, Y_k imply that the image of Y_k is a direct complement of the kernel of X_k and the image of X_k is a direct complement of the kernel of Y_k.

As $X_k Y_k = X_k E^{-1} E Y_k$ we can efficiently compute the product $X_k Y_k \in \mathrm{Adj}(\mathbf{B}')$. $X_k Y_k$ cannot be zero. Note that if $X_k Y_k$ is invertible, then X is also invertible, and the XE^{-1} in our hand sends $\mathbf{B}$ to $\mathbf{B}'$, which solves the problem. So in the following we assume $X_k Y_k$ has a non-trivial kernel.

Similarly to the stabilization argument above, we may assume that k is large enough so that the kernel of $X_k Y_k$ in the left regular representation $\mathrm{Adj}(\mathbf{B}')$ is a direct complement of the image. This mean that the right annihilator of $X_k Y_k$ in $\mathrm{Adj}(\mathbf{B}')$ (which is the same as that of X_k) and the right (as well as the left) ideal generated by $X_k Y_k$ (which is also generated by Y_k or X_k) are complementary to each other and the same holds for the product $Y_k X_k$.

We claim that there exists $\mathbf{B}'' \leq S^\epsilon(n, \mathbb{F})$ such that $\mathcal{B} = E'\mathbf{B}''$ for some invertible E' and $X_k Y_k \in L^\epsilon(\mathbf{B}'')$. To see this, consider an element $Z \in L^\epsilon(\mathbf{B}')$ which is a generator of the right annihilator of X_k as a right ideal in $\mathrm{Adj}(\mathbf{B}')$. Such Z exists by Lemma 4.1. Put $W = Y_k + Z$. Then $W \in L^\epsilon(\mathbf{B}')$, and W is invertible since Y_k and Z are generators of right ideals of $\mathrm{Adj}(\mathbf{B}')$ complementary to each other. We also have $X_k W = X_k(Y_k + Z) = X_k Y_k$. Let $\mathbf{B}'' = W^{-1}\mathbf{B}'$. Then, W^{-1} is an invertible element of $L^\epsilon(\mathbf{B}')$, so we have $\mathbf{B}'' \leq S^\epsilon(n, \mathbb{F})$. Furthermore, $L^\epsilon(\mathbf{B}'') = L^\epsilon(W^{-1}\mathbf{B}') = L^\epsilon(\mathbf{B}')W$. In particular, $X_k Y_k = X_k W \in L^\epsilon(\mathbf{B}')W = L^\epsilon(\mathbf{B}'')$.

Let J resp. K be the image resp. the kernel of $X_k Y_k$. From $X_k Y_k \in L^\epsilon(\mathbf{B}'')$ we infer $J = K^{\perp_{\mathbf{B}''}}$. Let $J' = K^{\perp_{\mathcal{B}}}$ and $K' = J^{\perp_{\mathcal{B}}}$. These subspaces can be computed efficiently. Let U_0 be an invertible linear map that maps J to J' and K to K'. Then by replacing $\mathcal{B}$ with $U_0^t \mathcal{B}$ we can arrange that $J = K^{\perp_{\mathcal{B}}}$ as well. Then the problem can be reduced to the subspaces J and K.

Acknowledgements. Part of this research was accomplished while the first author was visiting the Centre for Quantum Technologies, National University of Singapore. His research was also partially supported by the Hungarian National Research, Development and Innovation Office – NKFIH, Grant K115288. Y. Q.'s research was supported by the Australian Research Council DECRA DE150100720.

References

[Alb39] A.A. Albert. *Structure of Algebras*. Number v. 24 in American Mathematical Society colloquium publications. American Mathematical Society, 1939.

[Atk83] M. D. Atkinson. Primitive spaces of matrices of bounded rank. II. *Journal of the Australian Mathematical Society (Series A)*, 34(03):306–315, 1983.

[Bae38] Reinhold Baer. Groups with abelian central quotient group. *Trans. Amer. Math. Soc.*, 44:357–386, 1938.

[BFFP11] Charles Bouillaguet, Jean-Charles Faugère, Pierre-Alain Fouque, and Ludovic Perret. Practical cryptanalysis of the identification scheme based on the isomorphism of polynomial with one secret problem. In *Public Key Cryptography - PKC 2011 - 14th International Conference on Practice and Theory in Public Key Cryptography, Taormina, Italy, March 6-9, 2011. Proceedings*, pages 473–493, 2011.

[BFP15] Jérémy Berthomieu, Jean-Charles Faugère, and Ludovic Perret. Polynomial-time algorithms for quadratic isomorphism of polynomials: The regular case. *J. Complexity*, 31(4):590–616, 2015.

[BFV13] Charles Bouillaguet, Pierre-Alain Fouque, and Amandine Véber. Graph-theoretic algorithms for the "isomorphism of polynomials" problem. In *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, pages 211–227, 2013.

[BL08] Peter A. Brooksbank and Eugene M. Luks. Testing isomorphism of modules. *Journal of Algebra*, 320(11):4020 – 4029, 2008.

[BMW17] Peter A. Brooksbank, Joshua Maglione, and James B. Wilson. A fast isomorphism test for groups whose lie algebra has genus 2. *Journal of Algebra*, 473:545 – 590, 2017.

[BO08] Peter A Brooksbank and Eamonn A O'Brien. Constructing the group preserving a system of forms. *International Journal of Algebra and Computation*, 18(02):227–241, 2008.

[BW12] Peter A. Brooksbank and James B. Wilson. Computing isometry groups of Hermitian maps. *Trans. Amer. Math. Soc.*, 364:1975–1996, 2012.

[CIK97] Alexander Chistov, Gábor Ivanyos, and Marek Karpinski. Polynomial time algorithms for modules over finite dimensional algebras. In *Proceedings of the 1997 international symposium on Symbolic and algebraic computation*, ISSAC '97, pages 68–74, New York, NY, USA, 1997. ACM.

[CIKK15] Marco Carmosino, Russell Impagliazzo, Valentine Kabanets, and Antonina Kolokolova. Tighter connections between derandomization and circuit lower bounds. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2015, August 24-26, 2015, Princeton, NJ, USA*, pages 645–658, 2015.

[CJL+16] Lily Chen, Stephen Jordan, Yi-Kai Liu, Dustin Moody, Rene Peralta, Ray Perlner, and Daniel Smith-Tone. *Report on post-quantum cryptography.* US Department of Commerce, National Institute of Standards and Technology, 2016.

[dGIR96] Willem A. de Graaf, Gábor Ivanyos, and Lajos Rónyai. Computing cartan subalgebras of lie algebras. *Appl. Algebra Eng. Commun. Comput.*, 7(5):339–349, 1996.

[Dic23] Leonard Eugene Dickson. Algebras and their arithmetics. *University of Chicago Science Series*, 1923.

[Ebe91a] Wayne Eberly. Decomposition of algebras over finite fields and number fields. *Computational Complexity*, 1:183–210, 1991.

[Ebe91b] Wayne Eberly. Decompositions of algebras over R and C. *Computational Complexity*, 1:211–234, 1991.

[EG00] Wayne Eberly and Mark Giesbrecht. Efficient decomposition of associative algebras over finite fields. *Journal of Symbolic Computation*, 29(3):441–458, 2000.

[EH88] David Eisenbud and Joe Harris. Vector spaces of matrices of low rank. *Advances in Mathematics*, 70(2):135 – 155, 1988.

[FN70] V. Felsch and J. Neubüser. On a programme for the determination of the automorphism group of a finite group. In Pergamon J. Leech, editor, *Computational Problems in Abstract Algebra (Proceedings of a Conference on Computational Problems in Algebra, Oxford, 1967)*, pages 59–60, Oxford, 1970.

[FP06] Jean-Charles Faugère and Ludovic Perret. Polynomial equivalence problems: Algorithmic and theoretical aspects. In *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings*, pages 30–47, 2006.

[FR85] Katalin Friedl and Lajos Rónyai. Polynomial time solutions of some problems of computational algebra. In *Proceedings of the seventeenth annual ACM symposium on Theory of computing*, pages 153–162. ACM, 1985.

[GGOW16] Ankit Garg, Leonid Gurvits, Rafael Oliveira, and Avi Wigderson. A deterministic polynomial time algorithm for non-commutative rational identity testing. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pages 109–117, 2016.

[GMS03] Willi Geiselmann, Willi Meier, and Rainer Steinwandt. An attack on the isomorphisms of polynomials problem with one secret. *Int. J. Inf. Sec.*, 2(1):59–64, 2003.

[GQ17] Joshua A. Grochow and Youming Qiao. Algorithms for group isomorphism via group extensions and cohomology. *SIAM Journal on Computing*, 46(4):1153–1216, 2017.

[IKS10] Gábor Ivanyos, Marek Karpinski, and Nitin Saxena. Deterministic polynomial time algorithms for matrix completion problems. *SIAM J. Comput.*, 39(8):3736–3751, 2010.

[IQ17] Gábor Ivanyos and Youming Qiao. Algorithms based on *-algebras, and their applications to isomorphism of polynomials with one secret, group isomorphism, and polynomial identity testing. *CoRR*, abs/1708.03495, 2017.

[IQS17a] Gábor Ivanyos, Youming Qiao, and K. V. Subrahmanyam. Constructive non-commutative rank is in deterministic polynomial time. In *the 8th Innovations in Theoretical Computer Science (ITCS)*, pages 23:1–23:19, 2017.

[IQS17b] Gábor Ivanyos, Youming Qiao, and K. V. Subrahmanyam. Non-commutative edmonds' problem and matrix semi-invariants. *Computational Complexity*, 26(3):717–763, 2017.

[IR93] Gábor Ivanyos and Lajos Rónyai. Finding maximal orders in semisimple algebras over Q. *Computational Complexity*, 3:245–261, 1993.

[IRS12] Gábor Ivanyos, Lajos Rónyai, and Josef Schicho. Splitting full matrix algebras over algebraic number fields. *J. Algebra*, 354:211–223, 2012.

[Iva00] Gábor Ivanyos. Fast randomized algorithms for the structure of matrix algebras over finite fields. In *Proceedings of the 2000 international symposium on Symbolic and algebraic computation*, pages 175–183. ACM, 2000.

[Kay11] Neeraj Kayal. Efficient algorithms for some special cases of the polynomial equivalence problem. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 1409–1421, 2011.

[KI04] V. Kabanets and R. Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity*, 13(1/2):1–46, 2004.

[KL90] William M. Kantor and Eugene M. Luks. Computing in quotient groups. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, May 13-17, 1990, Baltimore, Maryland, USA*, pages 524–534, 1990.

[Lew77] DW Lewis. Forms over real algebras and the multisignature of a manifold. *Advances in Mathematics*, 23(3):272–284, 1977.

[Lew06] David W. Lewis. Involutions and anti-automorphisms of algebras. *Bulletin of the London Mathematical Society*, 38:529–545, 8 2006.

[Lov89] László Lovász. Singular spaces of matrices and their application in combinatorics. *Boletim da Sociedade Brasileira de Matemática-Bulletin/Brazilian Mathematical Society*, 20(1):87–99, 1989.

[LQ17] Yinan Li and Youming Qiao. Linear algebraic analogues of the graph isomorphism problem and the Erdős-Rényi model. In *FOCS*, pages 463–474, 2017.

[LW12] Mark L. Lewis and James B. Wilson. Isomorphism in expanding families of indistinguishable groups. *Groups - Complexity - Cryptology*, 4(1):73110, 2012.

[Mil78] Gary L. Miller. On the $n \log n$ isomorphism technique (a preliminary report). In *STOC*, pages 51–58, New York, NY, USA, 1978. ACM.

[MP08] Guillaume Malod and Natacha Portier. Characterizing valiant's algebraic complexity classes. *J. Complexity*, 24(1):16–38, 2008.

[MPG13] Gilles Macario-Rat, Jérôme Plût, and Henri Gilbert. New insight into the isomorphism of polynomial problem IP1S and its use in cryptography. In

Advances in Cryptology - ASIACRYPT 2013 - 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, India, December 1-5, 2013, Proceedings, Part I, pages 117–133, 2013.

[Pat96] Jacques Patarin. 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*, pages 33–48, 1996.

[PCDY17] Albrecht Petzoldt, Ming-Shing Chen, Jintai Ding, and Bo-Yin Yang. Hmfev - an efficient multivariate signature scheme. In *Post-Quantum Cryptography - 8th International Workshop, PQCrypto 2017, Utrecht, The Netherlands, June 26-28, 2017, Proceedings*, pages 205–223, 2017.

[Per05] Ludovic Perret. A fast cryptanalysis of the isomorphism of polynomials with one secret problem. In *Advances in Cryptology - EUROCRYPT 2005, 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Aarhus, Denmark, May 22-26, 2005, Proceedings*, pages 354–370, 2005.

[PFM14] Jérôme Plût, Pierre-Alain Fouque, and Gilles Macario-Rat. Solving the "isomorphism of polynomials with two secrets" problem for all pairs of quadratic forms. *CoRR*, abs/1406.3163, 2014.

[PGC98] Jacques Patarin, Louis Goubin, and Nicolas Courtois. Improved algorithms for isomorphisms of polynomials. In *Advances in Cryptology - EUROCRYPT '98, International Conference on the Theory and Application of Cryptographic Techniques, Espoo, Finland, May 31 - June 4, 1998, Proceeding*, pages 184–200, 1998.

[Pie82] Richard S. Pierce. *Associative algebras*, volume 88 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1982. Studies in the History of Modern Science, 9.

[Rón87] Lajos Rónyai. Simple algebras are difficult. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, pages 398–408, 1987.

[Rón90] Lajos Rónyai. Computing the structure of finite algebras. *Journal of Symbolic Computation*, 9(3):355–373, 1990.

[Rón92] Lajos Rónyai. Algorithmic properties of maximal orders in simple algebras over $\mathbf{Q}$. *Comput. Complexity*, 2(3):225–243, 1992.

[Rón94] Lajos Rónyai. A deterministic method for computing splitting elements in simple algebras over Q. *J. Algorithms*, 16(1):24–32, 1994.

[Sax09] Nitin Saxena. Progress on polynomial identity testing. *Bulletin of the EATCS*, 99:49–79, 2009.

[Sax13] Nitin Saxena. Progress on polynomial identity testing - II. *Electronic Colloquium on Computational Complexity (ECCC)*, 20:186, 2013.

[SY10] Amir Shpilka and Amir Yehudayoff. Arithmetic circuits: A survey of recent results and open questions. *Foundations and Trends in Theoretical Computer Science*, 5:207–388, March 2010.

[Taf57] E. J. Taft. Invariant wedderburn factors. *Illinois Journal of Mathematics*, 1(4):565–573, 1957.

[Tod92] S. Toda. Classes of arithmetic circuits capturing the complexity of computing the determinant. *IEICE Trans. Inf. Syst.*, E75-D:116–124, 1992.

[vdW05] Christiaan van de Woestijne. Deterministic equation solving over finite fields. In *Proceedings of the 2005 international symposium on Symbolic and algebraic computation*, pages 348–353. ACM, 2005.

[Wey97] H. Weyl. *The classical groups: their invariants and representations*, volume 1. Princeton University Press, 1997.

[Wil09a] James B. Wilson. Decomposing p-groups via Jordan algebras. *Journal of Algebra*, 322(8):2642–2679, 2009.

[Wil09b] James B. Wilson. Finding central decompositions of p-groups. *Journal of Group Theory*, 12(6):813–830, 2009.

[Wil09c] R. Wilson. *The Finite Simple Groups*, volume 251 of *Graduate Texts in Mathematics*. Springer London, 2009.

[Wil13] James B. Wilson. Optimal algorithms of gram–schmidt type. *Linear Algebra and its Applications*, 438(12):4573–4583, 2013.

[Wil14] James B. Wilson. 2014 conference on *Groups, Computation, and Geometry* at Colorado State University, co-organized by P. Brooksbank, A. Hulpke, T. Penttila, J. Wilson, and W. Kantor. Personal communication, 2014.

[Wol05] Christopher Wolf. Multivariate quadratic polynomials in public key cryptography. *IACR Cryptology ePrint Archive*, 2005:393, 2005.