

From independent sets and vertex colorings to isotropic spaces and isotropic decompositions: Another bridge between graphs and alternating matrix spaces*

Xiaohui Bei[†] Shiteng Chen[‡] Ji Guan[§] Youming Qiao[¶] Xiaoming Sun^{||}

December 30, 2020

Abstract

In the 1970's, Lovász built a bridge between graphs and alternating matrix spaces, in the context of perfect matchings (FCT 1979). A similar connection between bipartite graphs and matrix spaces plays a key role in the recent resolutions of the non-commutative rank problem (Garg-Gurvits-Oliveira-Wigderson, FOCS 2016; Ivanyos-Qiao-Subrahmanyam, ITCS 2017). In this paper, we lay the foundation for another bridge between graphs and alternating matrix spaces, in the context of independent sets and vertex colorings. The corresponding structures in alternating matrix spaces are *isotropic spaces* and *isotropic decompositions*, both useful structures in group theory and manifold theory.

We first show that the maximum independent set problem and the vertex c -coloring problem reduce to the maximum isotropic space problem and the isotropic c -decomposition problem, respectively. Next, we show that several topics and results about independent sets and vertex colorings have natural correspondences for isotropic spaces and decompositions. These include algorithmic problems, such as the maximum independent set problem for bipartite graphs, and exact exponential-time algorithms for the chromatic number, as well as mathematical questions, such as the number of maximal independent sets, and the relation between the maximum degree and the chromatic number. These connections lead to new interactions between graph theory and algebra. Some results have concrete applications to group theory and manifold theory, and we initiate a variant of these structures in the context of quantum information theory. Finally, we propose several open questions for further exploration.

Dedicated to the memory of Ker-I Ko

*A preliminary version of this paper was accepted to the 11th Innovations in Theoretical Computer Science conference (ITCS 2020) [BCG⁺20].

[†]School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore. Email address: xhbei@ntu.edu.sg.

[‡]State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences, Beijing, China. Email address: chenst@ios.ac.cn.

[§]State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences, Beijing, China. Email address: guanji1992@gmail.com.

[¶]Centre for Quantum Software and Information, University of Technology Sydney, Australia. Email address: jimmyqiao86@gmail.com.

^{||}Institute of Computing Technology, Chinese Academy of Sciences and University of Chinese Academy of Sciences, Beijing, China. Email address: sunxiaoming@ict.ac.cn.

Contents

1	Introduction	1
1.1	Between graphs and matrix spaces: some known bridges	1
1.2	Between graphs and matrix spaces: a new bridge	2
1.3	A summary of our results and their applications	4
2	A detailed description of the results	6
2.1	Analogues of bipartite testing and maximum independent set on bipartite graphs . .	6
2.2	Analogues of maximal independent sets and exact exponential-time algorithms of chromatic numbers	9
2.3	Complexity upper bound over \mathbb{C} , and independence number vs maximum degree . .	10
2.4	Applications of our results	11
2.5	Outlook	12
3	Preliminaries	13
3.1	Computational models	16
4	Basic facts and properties	16
4.1	Isotropic spaces and decompositions in group theory and manifold theory	18
5	Proof of Theorem 3	19
6	An exposition of the proof of Theorem 6	20
7	Proof of Theorem 8	22
8	Proof of Theorem 10	24
8.1	The existential singularity problem for matrix spaces	26
9	Proof of Theorem 11	28
9.1	Some basic statistics	28
9.2	A non-trivial upper bound on the number of maximal isotropic spaces	30
9.2.1	Proof of Lemma 34	33
9.3	Turning Theorem 11 into an algorithm	34
10	Proof of Theorem 12	35
11	Proofs for Propositions 13 and 14	37
12	Proofs of Theorems 15 and 16	39
13	A quantum variant of the theory	39
13.1	From connected graphs to irreducible quantum channels	40
13.2	Quantum gate subspace-fidelity and isotropic spaces	42

1 Introduction

1.1 Between graphs and matrix spaces: some known bridges

The bridge between perfect matchings and full-rank matrices. It is well-known that some graph-theoretic problems reduce to certain problems about linear spaces of matrices. A classical example, tracing back to Tutte [Tut47], and then more systematically examined by Lovász [Lov79, Lov89], concerns perfect matchings.

Let \mathbb{F} be a field, and $[n] := \{1, \dots, n\}$. Let $G = ([n], E)$ be a simple and undirected graph, so E can be viewed as a subset of $\{\{i, j\} : i, j \in [n], i \neq j\}$. For $n \in \mathbb{N}$ and $\{i, j\}$ where $i, j \in [n]$, $i < j$, the elementary alternating¹ matrix $A_{i,j}$ of size $n \times n$ is the matrix with the (i, j) th entry being 1, the (j, i) th entry being -1 , and the rest entries being 0. Let \mathcal{A}_G be the linear space of alternating matrices spanned by $A_{i,j}$, $\{i, j\} \in E$. Then when the field is large enough, G has a perfect matching if and only if \mathcal{A}_G contains a full-rank matrix.

A similar construction for bipartite graphs is also classical. Let $G = (L \cup R, E)$ be a bipartite graph where $L = R = [n]$, so E can be viewed as a subset of $[n] \times [n]$. For $n \in \mathbb{N}$ and $i, j \in [n]$, the elementary matrix $E_{i,j}$ of size $n \times n$ is the matrix with the (i, j) th entry being 1, and the rest entries being 0. Let \mathcal{B}_G be the linear space of matrices spanned by $E_{i,j}$, $(i, j) \in E$. Then when the field is large enough, G has a perfect matching if and only if \mathcal{B}_G contains a full-rank matrix.

As noted by Lovász [Lov79], these observations give efficient randomized algorithms for deciding the existence of perfect matchings on bipartite graphs and graphs over a large enough \mathbb{F} via the celebrated Schwartz-Zippel lemma [Zip79, Sch80]. Furthermore, because the determinant polynomial can be evaluated efficiently in parallel [Ber84, Chi85], these are actually randomized NC algorithms.

This work of Lovász has inspired several prominent results, including randomized NC algorithms for constructing perfect matchings [KUW86, MVV87], and the recent breakthrough of quasi-NC algorithms for perfect matchings on bipartite graphs [FGT16] and on general graphs [ST17]. Furthermore, derandomizing the corresponding algorithm for *general* linear spaces of matrices – not necessarily those of the form \mathcal{B}_G or \mathcal{A}_G – is now known as the symbolic determinant identity testing problem, and turns out to be of fundamental significance in complexity theory, as that would imply strong circuit lower bounds which are considered to be beyond current techniques [KI04, CIKK15].

In the following, we shall call linear spaces of (alternating) matrices as (alternating) matrix spaces. For a field \mathbb{F} , we use $M(s \times t, \mathbb{F})$ to denote the linear space of all $s \times t$ matrices over \mathbb{F} , and write $\mathcal{B} \leq M(s \times t, \mathbb{F})$ to denote that \mathcal{B} is a matrix space in $M(s \times t, \mathbb{F})$. Let $M(n, \mathbb{F}) := M(n \times n, \mathbb{F})$, and $\Lambda(n, \mathbb{F})$ be the linear space of all $n \times n$ alternating matrices over \mathbb{F} .

The bridge between shrunk subsets and shrunk subspaces. For bipartite graphs, a structure closely related to perfect matchings is the following. Given a bipartite graph $G = (L \cup R, E)$ where $L = R = [n]$, we say that a subset $S \subseteq L$ is a shrunk² subset of G , if $|S| > |N(S)|$ where $N(S)$ is the set of neighbours of S in R . The celebrated Hall’s marriage theorem [Hal35] says that G has a perfect matching if and only if it does not have a shrunk subset.

On the matrix space side, it is then natural to define the so-called shrunk subspaces. Specifically, given a matrix space $\mathcal{B} \leq M(n, \mathbb{F})$, a subspace $U \leq \mathbb{F}^n$ is a shrunk subspace of \mathcal{B} , if $\dim(U) > \dim(\mathcal{B}(U))$ where $\mathcal{B}(U) := \langle \cup B(U) : B \in \mathcal{B} \rangle$, and $B(U)$ denotes the image of U under B .

¹An $n \times n$ matrix A over \mathbb{F} is *alternating*, if for any $v \in \mathbb{F}^n$, $v^t A v = 0$. An alternating matrix is always skew-symmetric (i.e. $A^t = -A$), and a skew-symmetric matrix is also alternating over fields of characteristic not 2.

²We call S to be “shrunk” instead of “shrinking”, as we think of the bipartite graph G as shrinking the set S .

As in the perfect matching case, a bipartite graph G has a shrunk subset if and only if \mathcal{B}_G has a shrunk subspace [Lov89]. However, for general matrix spaces, the natural analogue of Hall’s theorem, namely a matrix space contains full-rank matrices if and only if it has no shrunk subspaces, does not hold, as evidenced by the space of all 3×3 alternating matrices. (The only if direction trivially holds, though.) Therefore, the bridge between shrunk subsets and shrunk subspaces is different from the one between perfect matchings and full-rank matrices.

The problem of testing whether a matrix space has a shrunk subspace arises naturally from several mathematical and computational disciplines, including algebraic complexity, non-commutative algebra, invariant theory, and analysis [GGdOW16, IQS18, GGdOW17]. Not surprisingly then, this problem has had several names. We adopt the *non-commutative rank problem* which seems widely used now, and refer an interested reader to [GGdOW16, IQS18] for the origin of this name.

With all these motivations, the non-commutative rank problem recently received considerable attention, and substantial progress has been made. First raised by Cohn [Coh75] four decades ago in the study of free fields, it was more recently reached at by Mulmuley in the context of derandomizing the Noether’s Normalization Lemma [Mul12, Mul17], and also by Hrubeš and Wigderson in the context of non-commutative arithmetic circuits with divisions [HW15]. Only known to be in PSPACE before 2015 [CR99], this problem was shown to be in P over the rational number field [GGdOW16] and over any field [IQS18].

The techniques supporting the solutions to the non-commutative rank problem are reminiscent of the corresponding techniques for the perfect matching problem on bipartite graphs. In [GGdOW16], it is the scaling algorithm [Sin64, LSW00], generalized to the quantum operator setting [Gur04]. In [IQS18], it is the classical augmenting path algorithm, generalized to the matrix space setting [IKQS15, IQS17]. Ingredients from invariant theory are also crucial. For [GGdOW16], Garg et al. needed the exponential upper bound on generating the ring of matrix semi-invariants [Der01]. For [IQS18], Ivanyos et al. need the polynomial upper bound [DM17], which in turn relies crucially on the regularity lemma developed in [IQS17].

1.2 Between graphs and matrix spaces: a new bridge

In this paper we lay the foundation for yet another bridge between graphs and matrix spaces. We focus on undirected simple graphs, hence it is natural, as Tutte and Lovász did with perfect matchings, to work with alternating matrix spaces. We start from independent sets and vertex colorings, two central structures in graph theory with numerous results from various motivations [JT95, Die17]. By identifying analogues of them in the alternating matrix space setting, we arrive at isotropic spaces and isotropic decompositions, which we define now.

Definition 1. *Let $\mathcal{A} \leq \Lambda(n, \mathbb{F})$ be an alternating matrix space. A subspace $U \leq \mathbb{F}^n$ is an isotropic space of \mathcal{A} , if for any $u, u' \in U$, and any $A \in \mathcal{A}$, we have $u^t A u' = 0$. For $c \in \mathbb{N}$, an isotropic c -decomposition of \mathcal{A} is a direct sum decomposition of \mathbb{F}^n into c non-zero subspaces $U_1 \oplus \dots \oplus U_c$, where every U_i is an isotropic space.*

Recall that for a graph $G = ([n], E)$, an independent set of G is a subset $S \subseteq [n]$ such that for any $i, j \in S$, there is no edge from E connecting these two vertices. A vertex c -coloring of G is a partition of the vertex set into c independent sets. Therefore, the definitions of isotropic spaces and isotropic decompositions do mimic those of independent sets and vertex colorings. It is then natural to introduce the following definitions and the corresponding algorithmic problems.

Definition 2. Let $\mathcal{A} \leq \Lambda(n, \mathbb{F})$. The isotropic number of \mathcal{A} , denoted as $\alpha(\mathcal{A})$, is the maximum $d \in \mathbb{N}$ such that \mathcal{A} has an isotropic space of dimension d . The isotropic decomposition number, denoted as $\chi(\mathcal{A})$, is the minimum $c \in \mathbb{N}$ such that \mathcal{A} admits an isotropic c -decomposition.

Given $d \in \mathbb{N}$ and a linear basis of $\mathcal{A} \leq \Lambda(n, \mathbb{F})$, the maximum isotropic space problem asks to decide whether $\alpha(\mathcal{A}) \geq d$. Given $c \in \mathbb{N}$ and a linear basis of $\mathcal{A} \leq \Lambda(n, \mathbb{F})$, the isotropic c -decomposition problem asks to decide whether $\chi(\mathcal{A}) \leq c$.

Note that $\alpha(\cdot)$ and $\chi(\cdot)$ are used to denote the independent number and the chromatic number of graphs [Die17], and these choices are deliberate. Also note that for any $\mathcal{A} \leq \Lambda(n, \mathbb{F})$, we have $\alpha(\mathcal{A}) \geq 1$, and $\chi(\mathcal{A}) \leq n$. Indeed, due to the alternating condition, any $\mathcal{A} \leq \Lambda(n, \mathbb{F})$ enjoys the property that any 1-dimensional subspace of \mathbb{F}^n is an isotropic space of \mathcal{A} . It follows that any direct sum decomposition of \mathbb{F}^n into n dimension-1 subspaces is an isotropic n -decomposition of \mathcal{A} . This property corresponds nicely to that for any undirected simple graph, every single vertex is an independent set. On the other hand, symmetric matrix spaces do not satisfy this property in general. Therefore, this small but pleasant coincidence suggests that working with alternating matrix spaces is a natural choice in this setting.

Our first result follows what Lovász did with perfect matchings, and provides a first indication on the new connection. Recall that given a graph $G = ([n], E)$, we can associate an alternating matrix spaces $\mathcal{A}_G \leq \Lambda(n, \mathbb{F})$, spanned by those elementary alternating matrices $A_{i,j}$ with $\{i, j\} \in E$.

Theorem 3. Let G and \mathcal{A}_G be as above. Then we have

1. G has a size- s independent set if and only if \mathcal{A}_G has a dimension- s isotropic space. In particular, $\alpha(G) = \alpha(\mathcal{A}_G)$.
2. G has a vertex c -coloring if and only if \mathcal{A}_G has an isotropic c -decomposition. In particular, $\chi(G) = \chi(\mathcal{A}_G)$.

The proof is in Section 5. Theorem 3 demonstrates that the maximum isotropic space problem and the isotropic decomposition problem are genuine generalizations³ of the maximum independent set problem and the vertex c -coloring problem, respectively. It also implies the following.

Corollary 4. The maximum isotropic space problem and the isotropic 3-decomposition problem for alternating matrix spaces are NP-hard.

Emboldened by Theorem 3, we propose to view isotropic spaces and decompositions as linear algebraic analogues of independent sets and vertex colorings, and study these two structures from the perspectives of graph theory and algorithms. This leads to natural and interesting mathematical and algorithmic problems, whose solutions bring together strategies, techniques, and results from several areas, including graph theory, algorithm design, computer algebra, and algebraic complexity. We regard these results as laying the foundation of a new bridge between graphs and alternating matrix spaces.

While our investigation started with an analogy, isotropic spaces and decompositions are actually classical notions, with natural interpretations in group theory and manifold theory. Therefore, some of our results have concrete applications to these areas. We also initiate a variant of this theory

³Note that for the maximum isotropic space problem and the isotropic decomposition problem, we consider all alternating matrix spaces, not necessarily of the form \mathcal{A}_G coming from a graph G .

in quantum information, and find an interesting information theoretic interpretation of isotropic spaces in quantum error correction. These demonstrate the usefulness of this new bridge.

To describe our results, we set up some notation. We use \mathbb{F}_q , \mathbb{Q} , \mathbb{R} , and \mathbb{C} to denote the finite field with q elements, the rational number field, the real number field, and the complex number field, respectively. Elements in \mathbb{F}^n are *column* vectors. In algorithms, subspaces of \mathbb{F}^n and $\Lambda(n, \mathbb{F})$ are represented by linear bases. We may write $\Lambda(n, \mathbb{F}_q)$ as $\Lambda(n, q)$ for convenience. For more details on the computation model, see Section 3.1.

1.3 A summary of our results and their applications

We would like to provide a summary of our results and their applications before a detailed description in Section 2.

Because independent sets and vertex colorings are classical and central topics in graph theory, many problems and techniques are studied and developed centred around these two notions. Our new bridge then allows for transferring a host of problems from graphs to alternating matrix spaces. For example, let us recall some well-known graph-theoretic and/or algorithmic results on independent sets and vertex colorings (see e.g. [Die17]).

1. Whether a graph is bipartite can be tested in deterministic polynomial time.
2. On bipartite graphs, the maximum independent set problem is in P.
3. Any n -vertex and m -edge graph has an independent set of size $\geq \frac{n^2}{2m+n}$ [Tur41]⁴.
4. The number of maximal independent sets on an n -vertex graph is $\leq 3^{\frac{n}{3}}$ [MM65].
5. The chromatic number of an n -vertex graph can be computed in time $(1 + 3^{\frac{1}{3}})^n \cdot \text{poly}(n)$ [Law76]⁵.
6. The chromatic number of any graph is upper bounded by the maximum degree plus 1.

All of the above graph-theoretic results, and several more, will be found to have natural correspondences in the alternating matrix space setting.

With the abundance of results comes the difficulty of introducing and organising them. Indeed, it seems necessary to cover certain background information for each result from the above list, even though such information is well-known to most readers with an adequate knowledge in graph theory. Furthermore, the solutions to some results are not difficult, or are basically already in the literature, due to the connection between alternating matrix spaces and p -groups, and the connection between perfect matchings (resp. shrunk subsets) and full-rank matrices (resp. shrunk subspaces).

Still, despite the above concerns, we feel that the new bridge reveals somewhat unexpected links among different research directions and fields, and it is appropriate to introduce these results properly. So we shall take a detailed approach in Section 2 to go through all the results. Before that, for readers' convenience and for ease of comparison with the graph theory setting, let us also offer a summary of most results to be presented.

In Section 2.1, the following problems will be studied.

⁴This follows from Turán's celebrated result in extremal graph theory, which is usually stated for cliques, and implies this by simply taking the complement graph.

⁵This classical result of Lawler was from the 1970's, and the current status of the art is $2^n \cdot \text{poly}(n)$ [BHK09].

1. Bipartite testing of graphs corresponds to testing whether an alternating matrix space admits an isotropic 2-decomposition. The latter problem turns out to be in BPP over finite fields of odd size [BMW17] and in P over \mathbb{R} and \mathbb{C} [IQ19] by existing results (Theorem 6). The technique, based on $*$ -algebra approach developed by J. B. Wilson [Wil09a, Wil09b], is much more involved than the graph theory counterpart.
2. The maximum isotropic space problem for “bipartite” alternating matrix spaces turns out to be in BPP over finite fields of odd size. The key observation is a connection between the isotropic number of such an alternating matrix space, and the non-commutative rank of its “bipartite component” (Theorem 8). Given this observation, the maximum isotropic space problem in this case is in BPP (Corollary 9), by using two big hammers, namely the algorithm from [BMW17] as mentioned above, and the efficient algorithm for the non-commutative rank problem over finite fields [IQS17, IQS18].
3. It turns out that the correspondence of Turán’s extremal graph theory result in [Tur41] has been studied in group theory [Ol’78, BGH87], in the context of abelian subgroups of p -groups. It is then natural to compare the quantities to see some interesting phenomena.

In Section 2.1, we also study the problem of deciding whether an alternating matrix space admits a 2-dimensional isotropic space. While the corresponding problem is trivial on graphs, this problem turns out to be not so easy over \mathbb{Q} due to a reduction from quadratic residuosity problem (Theorem 10), which heavily relies on a result of Rónyai [Rón87]. One ingredient there is the introduction of the existential singularity problem for matrix spaces, which turns out to have rich connections to several mathematical disciplines (see Problem 1 and Section 8.1).

In Section 2.2, the following problems will be studied. *We view results in this direction as our main concrete contribution, after the key conceptual contribution of establishing a new bridge.*

4. We study the number of maximal isotropic spaces of alternating matrix spaces in $\Lambda(n, \mathbb{F}_q)$. Improving the trivial bound $q^{\frac{1}{4}n^2+O(n)}$, we prove a non-trivial upper bound $q^{\frac{1}{6}n^2+O(n)}$ (Theorem 11), and observe an easy lower bound $q^{\frac{1}{8}n^2+\Omega(n)}$ (Proposition 31).

We believe that the technique for proving the upper bound is worth noting. We adapt the proof strategy of the upper bound on maximal independent sets by Wood [Woo11]. This requires analogues of certain graph-theoretic concepts such as degrees and neighbours in the alternating matrix space setting, which have been developed in [Qia19]. It works up to some point, but after that, we have to resort to certain linear algebraic techniques. We leave closing the gap between $q^{\frac{1}{8}n^2+\Omega(n)}$ and $q^{\frac{1}{6}n^2+O(n)}$ an interesting open problem.

This result has a natural interpretation in group theory, as maximal isotropic subspaces of alternating matrix spaces over \mathbb{F}_p correspond to maximal abelian subgroups of p -groups of class 2 and exponent p for odd p .

5. We show that the isotropic decomposition number can be computed in time $q^{\frac{5}{12}n^2+O(n)}$ (Theorem 12), by a combination of enumerating maximal isotropic spaces with the dynamic programming technique. This is in close analogy with Lawler’s exact exponential-time algorithm for computing the chromatic number [Law76].

In Section 2.3, the following problems will be studied.

6. We study the relation between the isotropic decomposition number and a certain maximum degree notion of alternating matrix spaces. Such a degree notion was already useful in upper bounding the number of maximal isotropic spaces as above. This problem is also motivated by several works in algebra [Sch76, GL84, GG07].

Unlike in the graph setting, where the maximum degree plus one upper bounds the chromatic number, here we are only able to prove a weaker, but hopefully still interesting, bound in Proposition 14. Though not difficult, we believe that it leads to interesting problems for further investigation.

In Section 2.3, we also study the complexity-theoretic upper bounds for the maximum isotropic space problem and the isotropic 3-decomposition problem over \mathbb{C} in Proposition 13. The key observation is to use the celebrated result of Koiran on the Hilbert Nullstellensatz problem [Koi96].

After studying problems on alternating matrix spaces mostly by way of analogy with graphs, it is natural to ask whether some results have concrete applications. The answer is quite affirmative.

In Section 2.4, we shall provide two applications to *group theory*, one being computational, and the other being enumerative. As will be explained there, these applications are based on p -groups of class 2 and exponent p for odd p , for which testing isomorphism has long been known to be difficult, and is becoming more urgent in light of Babai's recent breakthrough on graph isomorphism [Bab16]. Briefly speaking, the computational problem is to compute the maximum abelian subgroup in the model of matrix groups over finite fields, and the enumerative problem is to bound the number of maximal abelian subgroups for p -groups of class 2 and exponent p .

We will then describe a variant of our theory in the context of *quantum information theory*, and present an information theoretic interpretation of isotropic spaces in the context of quantum error correction. In Section 4.1, we also present the connections to *manifold theory*, and mention a potential application. All these suggest that our results could be of interest to group theorists, quantum information theorists, and geometers, in particular to those who work on the computational aspects of these disciplines.

In Section 2.5, we shall summarize our contributions, and more importantly, present several open problems for further study.

The structure of the paper. We present a detailed description of all results in Section 2. We then present certain preliminaries in Section 3. Then in Section 4, we collect some basic facts and properties about isotropic spaces and decompositions, including their meanings in group theory and manifold geometry in Section 4.1. We then prove Theorem 3 in Section 5, which is the basis connecting the graph-theoretic structures and those structures on alternating matrix spaces. We then prove all the main results mentioned above in the following sections. (We will mention the corresponding section numbers when describing those results in Section 2.)

2 A detailed description of the results

2.1 Analogues of bipartite testing and maximum independent set on bipartite graphs

After Corollary 4, the isotropic 2-decomposition problem is of particular interest, as the vertex 2-coloring problem just asks whether a graph is bipartite, which can be tested efficiently by breadth-

first search. A moment's thought suggests that a breadth-first search type idea seems not applicable to the isotropic 2-decomposition problem.

Remark 5. Suppose $\mathcal{A} \leq \Lambda(n, \mathbb{F})$ admits an isotropic 2-decomposition as $\mathbb{F}^n = U_1 \oplus U_2$. Note that U_1 and U_2 are not known to us. To follow the idea of breadth-first search, we would start from a vector $v \in \mathbb{F}^n$, and then find its neighbours, and then its neighbours' neighbours, etc.. Intuitively, for $v \in \mathbb{F}^n$, we can view the linear span of Av , $A \in \mathcal{A}$ as those neighbours of v , denoted by $V_1 \leq \mathbb{F}^n$. Then the linear span of AV_1 , $A \in \mathcal{A}$, may be considered as the neighbours of V_1 . Continuing this way, if $v \in U_1$, we do see that V_i 's alternates between subspaces of U_1 and U_2 . It follows that $V_i \cap V_{i+1} = \mathbf{0}$, from which we can compute U_1 and U_2 after this sequence stabilizes. However, if v is neither in U_1 nor in U_2 , it is not clear how to read any information about U_1 and U_2 . In fact, it is possible that the linear span of Av is the hyperplane orthogonal to v , so it is impossible to tell whether such U_1 and U_2 exist.

Fortunately, it turns out that this problem has been studied in computer algebra over finite fields by Brooksbank, Maglione, and Wilson in [BMW17]. Their strategy can be readily applied to \mathbb{R} and \mathbb{C} , using some ingredients from [IQ19].

Theorem 6 ([BMW17, Theorem 3.6], [IQ19]). *The isotropic 2-decomposition problem can be solved in randomized polynomial time over \mathbb{F}_q with q odd, and in deterministic polynomial time over \mathbb{R} and \mathbb{C} . Furthermore, over \mathbb{F}_q with q odd, the algorithm also outputs the linear bases of the two subspaces in an isotropic 2-decomposition.*

While a proof for \mathbb{F}_q was already sketched in [BMW17], we still give an exposition of this proof in Section 6. Besides indicating how to handle \mathbb{R} and \mathbb{C} , we wish to give some reader a flavor of how the so-called *-algebra technique, pioneered by Wilson [Wil09a, Wil09b] in computer algebra, is applied to this setting. This technique was recently shown to be useful in polynomial identity testing and multivariate cryptography [IQ19].

Theorem 6 and its proof reveal that the isotropic spaces and the isotropic decompositions do have connections with, and implications to, other disciplines, just like the case of non-commutative ranks. Furthermore, a quantum variant of the theory can be developed, and the corresponding isotropic 2-decomposition problem can be solved efficiently using quantum information theoretic techniques (see Section 13). As mentioned above, the techniques used to solve the non-commutative rank problem also have their roots in algebra [IQS18] and quantum information [GGdOW16]. Perhaps it is not so coincidental that techniques from these areas are useful again.

In fact, the non-commutative rank problem arises naturally in our context. Since a bipartite graph is just a graph admitting a vertex 2-coloring, it is natural to make the following definition.

Definition 7. *An alternating matrix space is bipartite, if it admits an isotropic 2-decomposition.*

A well-known fact in graph theory is that, on bipartite graphs, the maximum independent set problem can be solved in deterministic polynomial time, through a reduction to the minimum vertex cover problem. The latter problem is equivalent to the maximum matching problem via König's theorem. It is then interesting to examine whether bipartite alternating matrix spaces admit an efficient algorithm for the maximum isotropic space problem. It turns out that the isotropic number of a bipartite alternating matrix space is closely related to the non-commutative rank of some matrix space, as we shall see now.

We have mentioned the decision version of the non-commutative rank problem in Section 1.1. We now define the non-commutative rank in a slightly more general setting. Given $\mathcal{B} \leq M(s \times t, \mathbb{F})$, its non-commutative rank is $\text{ncrk}(\mathcal{B}) := s+t - \max\{\dim(U) + \dim(V) : \forall u \in U, v \in V, B \in \mathcal{B}, u^t B v = 0\}$ [FR04]. Note that the recent works [GGdOW16, IQS18] used a slightly different formulation in the setting $s = t$.

Given a bipartite $\mathcal{A} \leq \Lambda(n, \mathbb{F})$, up to isometry (i.e. the action of $T \in \text{GL}(n, \mathbb{F})$ sending \mathcal{A} to $T^t \mathcal{A} T := \{T^t A T : A \in \mathcal{A}\}$), every $A \in \mathcal{A}$ is of the form $\begin{bmatrix} \mathbf{0} & B \\ -B^t & \mathbf{0} \end{bmatrix}$, where B is of size $s \times t$ (see Section 4). Let $\mathcal{B} \leq M(s \times t, \mathbb{F})$ be the space of such B arising from some $A \in \mathcal{A}$. Then we have:

Theorem 8. *Let $\mathcal{A} \leq \Lambda(n, \mathbb{F})$ and $\mathcal{B} \leq M(s \times t, \mathbb{F})$ be from above. Then $\alpha(\mathcal{A}) = n - \text{ncrk}(\mathcal{B})$.*

Thanks to the solution of the non-commutative rank problem over any field [IQS18], and Theorem 6 in the case of \mathbb{F}_q with odd q , we have

Corollary 9. *The maximum isotropic space problem for bipartite alternating matrix spaces of size $n \times n$ over \mathbb{F}_q , q odd, can be solved in randomized poly($n, \log q$) time.*

The proofs of Theorem 8 and Corollary 9 are in Section 7. In some sense, the non-commutative rank may be considered as corresponding to the minimum vertex cover size in the bipartite alternating matrix space setting. However, unlike in the graph case, where the relation between independent sets and vertex covers is so straightforward, the proof of Theorem 8 requires some twists, because of the “flexibility” of vectors and matrices.

Having seen the implication of the non-commutative rank problem to our setting, let us examine the following mathematical problem that arises naturally in our context, whose solution turns out to come from algebraic geometry. Again, let us trace back to the graph setting, and consider $\alpha(n, m) := \min\{\alpha(G) : G \text{ a graph with } n \text{ vertices and } m \text{ edges}\}$, where $\alpha(G)$ is the independence number. A celebrated result of Turán [Tur41] in extremal graph theory implies that

$$\alpha(n, m) \leq \left\lceil \frac{n^2}{2m + n} \right\rceil. \quad (1)$$

Turning to the alternating matrix space setting, it is natural to define $\alpha(\mathbb{F}, n, m) := \min\{\alpha(\mathcal{A}) : \mathcal{A} \leq \Lambda(n, \mathbb{F}), \dim(\mathcal{A}) = m\}$. This quantity has been studied by Buhler, Gupta, and Harris [BGH87] in relation to abelian subgroups of p -groups [Bur13, Alp65b]. The main result of [BGH87], proved using algebraic geometric techniques, is as follows: for any $m > 1$, we have

$$\alpha(\mathbb{F}, n, m) \leq \left\lfloor \frac{m + 2n}{m + 2} \right\rfloor, \quad (2)$$

where the equality is attainable over algebraically closed fields⁶. This inequality was also obtained earlier by Ol’shanskii [Ol’78]. Comparing Equations 1 and 2, we see that $\alpha(n, m)$ and $\alpha(\mathbb{F}, n, m)$ behave quite differently. For example, by Equation 1, every graph with n vertices and $2n$ edges has an independent set of size at least $n/5$. On the other hand, by Equation 2, there exists a dimension- $2n$ alternating matrix space in $\Lambda(n, \mathbb{F})$ with no isotropic space of dimension ≥ 2 .

Motivated by the discussion in the last paragraph, we study the algorithmic problem of deciding whether there exists an isotropic space of dimension ≥ 2 for $\mathcal{A} \leq \Lambda(n, \mathbb{F})$. This is equivalent to ask

⁶While in [BGH87] the main result was stated for fields of characteristic $\neq 2$, the proof, at least the inequality, works for any characteristic.

whether \mathcal{A} has an isotropic $(n-1)$ -decomposition. Note that the corresponding problem on graphs is trivial, as a graph has an independent set of size ≥ 2 if and only if it is not the complete graph. It turns out that over \mathbb{Q} , this problem is substantially more difficult.

Theorem 10. *Over \mathbb{Q} , assuming the generalized Riemann hypothesis, there is a randomized polynomial-time reduction from deciding quadratic residuosity modulo squarefree composite numbers to the problem of deciding whether an alternating matrix space has an isotropic space of dimension ≥ 2 .*

The proof is in Section 8, and it relies crucially on Rónyai’s fundamental work on computing algebra structures [Rón87].

2.2 Analogues of maximal independent sets and exact exponential-time algorithms of chromatic numbers

An independent set on a graph is maximal if it is not properly contained in some other independent set. The study of maximal independent sets is a classical demonstration of how graph theory and algorithm study are intertwined.

In the 1960’s, Erdős and Moser raised the question of bounding the number of maximal independent sets on a graph. It was subsequently solved by Moon and Moser [MM65], and alternative proofs have been found [Woo11, Vat11]. They show that the number of maximal independent set of an n -vertex graph is upper bounded by $3^{n/3}$, and this bound is tight. (Some refinement is required when n is not a multiple of 3.) Since the 1970’s, the problem of outputting all maximal independent sets or maximal cliques received considerable attention [AM70, TIAS77, LLK80, JPY88]. One application was provided by Lawler [Law76], who showed that the Moon-Moser bound together with dynamic programming give an algorithm for computing the chromatic number of an n -vertex graph in time $(1 + \sqrt[3]{3})^n \cdot \text{poly}(n)$. This algorithm was the starting point of exact exponential-time algorithms for chromatic numbers. Subsequent improvements [Epp03, Bys04, BH08] lead to the breakthrough by Björklund, Husfeldt, and Koivisto, who presented an algorithm in time $2^n \cdot \text{poly}(n)$ [BHK09].

Getting back to alternating matrix spaces, the natural correspondence would be maximal isotropic spaces. Formally, for an alternating matrix space $\mathcal{A} \leq \Lambda(n, \mathbb{F})$, an isotropic space is *maximal*, if there is no isotropic space properly containing it. We then ask analogous questions over finite fields, namely upper bounding the number of maximal isotropic spaces of $\mathcal{A} \leq \Lambda(n, \mathbb{F}_q)$, and exact exponential-time algorithms for computing the isotropic decomposition number $\chi(\mathcal{A})$. Interestingly, on one hand, these problems demonstrate behaviours different from the combinatorial counterpart. On the other hand, strategies and techniques from graph theory and algorithm design do carry over, in a non-trivial way, to these problems. Again, such phenomena have been witnessed in the non-commutative rank problem, and it is interesting to see these happening in this context. Furthermore, our result on the number of maximal isotropic spaces has a direct application to group theory, as we will see in Section 2.4.

We now describe our results in more details. To start with, we note that, as in the graph setting, an easy greedy algorithm outputs one maximal isotropic space (see Proposition 20). We then consider the number of maximal isotropic spaces for alternating matrix spaces in $\Lambda(n, q)$, analogously as done by Moon and Moser for graphs [MM65]. A trivial upper bound is the number of all subspaces of \mathbb{F}_q^n . This number, $q^{\frac{1}{4}n^2 + \Theta(n)}$, is well-known and classical (see Fact 30). Any alternating matrix space spanned by a single full-rank alternating matrix has $q^{\frac{1}{8}n^2 + \Theta(n)}$ maximal

isotropic spaces, providing a lower bound. This is also classical but perhaps not that well-known (see Proposition 31). We show a non-trivial upper bound as follows.

Theorem 11. *The number of maximal isotropic spaces of any $\mathcal{A} \leq \Lambda(n, q)$ is upper bounded by $q^{\frac{1}{6}n^2 + O(n)}$.*

The proof is in Section 9. As mentioned in Section 1.3, its proof strategy is inspired by Wood’s proof of the upper bound on maximal independent sets [Woo11]. This proof is constructive (see Corollary 36), so we can enumerate all maximal isotropic spaces in time $q^{\frac{1}{6}n^2 + O(n)}$. We then consider the problem of computing the isotropic decomposition number for $\mathcal{A} \leq \Lambda(n, q)$. A naive brute-force algorithm, namely enumerating all direct sum decompositions of \mathbb{F}_q^n , runs in time $q^{n^2 + O(n)}$. Inspired by Lawler’s strategy in [Law76], we combine our Corollary 36 with a dynamic programming idea to obtain the following.

Theorem 12. *The isotropic decomposition number of $\mathcal{A} \leq \Lambda(n, q)$ can be computed in time $q^{\frac{5}{12}n^2 + O(n)}$.*

The proof is in Section 10. An open question is whether the strategy in [BHK09] for chromatic numbers can be adapted to obtain an algorithm for isotropic decomposition numbers in time $q^{\frac{1}{4}n^2 + O(n)}$. This is because the number of subspaces of \mathbb{F}_q^n is $q^{\frac{1}{4}n^2 + \Theta(n)}$, while the algorithm in [BHK09] runs in time $2^n \cdot \text{poly}(n)$ where 2^n is the number of subsets of $[n]$.

2.3 Complexity upper bound over \mathbb{C} , and independence number vs maximum degree

We first consider complexity-theoretic upper bounds for the maximum isotropic space problem and the isotropic 3-decomposition problem. Clearly, these problems are in NP over finite fields. Over \mathbb{C} , we have the following, by resorting to a celebrated result of Koiran on the Hilbert Nullstellensatz problem [Koi96]. The proof is in Section 11.

Proposition 13. *Let $\mathcal{A} \leq \Lambda(n, \mathbb{C})$ be given by a linear basis consisting of integral matrices, and let c be a constant. The maximum isotropic space problem and the isotropic c -decomposition problem are in PSPACE unconditionally, and in PH assuming the generalized Riemann hypothesis.*

Our next result has two diverse motivations.

The first motivation is from linear algebra. Given a single alternating matrix A , its canonical form suggests that $\langle A \rangle$ admits an isotropic 2-decomposition (see Section 3). Given a pair of alternating matrices A_1 and A_2 , it is also known that $\langle A_1, A_2 \rangle$ admits an isotropic 2-decomposition [Sch76, GL84, GG07] (see also [BMW17, Lemma 3.7]). A natural question is what happens for alternating matrix spaces of dimension 3, or in general, any constant c .

The second motivation is from graph theory. Given a graph $G = ([n], E)$, let $\Delta(G)$ be the maximum degree over vertices of G . It is well-known that a simple greedy algorithm yields that $\chi(G) \leq \Delta(G) + 1$ [Die17, pp. 122]. For $\mathcal{A} \leq \Lambda(n, \mathbb{F})$, the *degree* of $v \in \mathbb{F}^n$ in \mathcal{A} can be defined as $\deg_{\mathcal{A}}(v) := \dim(\langle Av : A \in \mathcal{A} \rangle)$ [Qia19]. As mentioned in Section 2.2, this notion was already useful in the proof of Theorem 11. Let $\Delta(\mathcal{A}) = \max\{\deg_{\mathcal{A}}(v) : v \in \mathbb{F}^n\}$. It is then natural to ask the relation between $\chi(\mathcal{A})$ and $\Delta(\mathcal{A})$ in analogy to the graph setting. This question is closely related to the one in the last paragraph, since $\deg_{\mathcal{A}}(v) \leq \dim(\mathcal{A})$ for any $v \in \mathbb{F}^n$, so $\Delta(\mathcal{A}) \leq \dim(\mathcal{A})$.

We now present the following result, also deduced from a greedy algorithm.

Proposition 14. *Let $\mathcal{A} \leq \Lambda(n, \mathbb{F})$. Then $\chi(\mathcal{A}) \leq O(\Delta(\mathcal{A}) \cdot \log n)$. Furthermore, an isotropic C -decomposition with $C = O(\Delta(\mathcal{A}) \cdot \log n)$ can be found in polynomial time.*

The proof is in Section 11. Note that this implies that when $\dim(\mathcal{A})$ is a constant, then $\chi(\mathcal{A}) \leq O(\log n)$. We leave it an open problem for further improvement of the bound in Proposition 14.

2.4 Applications of our results

To start with, we recognize that the applications and connections to be presented are not surprising to some readers, because alternating bilinear maps, and therefore alternating matrix tuples and spaces⁷, naturally arise in group theory via the commutator bracket, and in manifold geometry via the cup product in cohomology. Therefore, certain isotropic spaces and decompositions have natural group-theoretic or geometric interpretations (see Section 4.1).

On the other hand, these applications may look exotic to some other readers, as they will be stated purely in group theoretic or quantum information theoretic terms. This is natural and expected, after a bridge is built. Indeed, the present bridge enables us to transfer problems, techniques, and results in graph theory and algorithm study, to other mathematical and computational disciplines which otherwise seem barely related to graph theory.

We now describe the first application to finite groups, more specifically, to computing with matrix groups over finite fields. Matrix groups over finite fields given by generators form an important model of computing with finite groups. In theoretical computer science, the study of this model led to the inventions of the black-box group model by Babai and Szemerédi [BS84], and the Arthur-Merlin class by Babai [Bab85]. Though some algorithms with worst-case analyses can be found in [Bea95], even the very basic membership testing problem was only recently known to be solvable in randomized polynomial time under a number-theoretic oracle [BBS09].

Overall, our knowledge about this model is rather limited, and many questions await investigations. One interesting problem is to compute an abelian subgroup of the largest size. Large abelian subgroups, besides motivations from computational group theory [Ros11], are useful in controlling the character degrees of the group, which in turn are useful in the group-theoretic approach for fast matrix multiplication [CKSU05]. As a consequence of Corollary 4, we have the following result, whose proof is in Section 12.

Theorem 15. *Let p be an odd prime. Given a matrix group G over \mathbb{F}_p , and $s \in \mathbb{N}$, deciding whether G has an abelian subgroup of order $\geq s$ is NP-hard.*

The proof of Theorem 15 relies on the connection between alternating matrix spaces over \mathbb{F}_p , and p -groups of class 2 and exponent p for odd p , via Baer’s correspondence [Bae38] (see Section 4.1).

It has long been known that p -groups of class 2 and exponent p form a bottleneck for testing isomorphism of finite groups. To solve the group isomorphism problem in time polynomial in the group order is a long-standing open problem [Mil78]. This problem is becoming more prominent in light of Babai’s breakthrough on graph isomorphism [Bab16], as Babai indicated the group isomorphism problem as a key bottleneck to put graph isomorphism in P [Bab16, arXiv version, Section 13.2]. Some interesting progress on testing isomorphism of such p -groups was recently made by utilizing the connection to alternating matrix spaces [LQ17].

Let us turn to the second application to finite groups. The question of bounding the number of maximal abelian subgroups has been considered for various group families [Dix71, Woo89, Ant87,

⁷For the relations among alternating matrix tuples and spaces, and alternating bilinear maps, see Sec. 3.

Vdo04], but to the best of our knowledge, there had been no results on this question for p -groups of class 2 and exponent p . Let P be such a group, so that the center $Z(P) \cong \mathbb{Z}_p^m$ and the central quotient $P/Z(P) \cong \mathbb{Z}_p^n$. The number of maximal abelian subgroups is upper bounded trivially by $p^{\frac{1}{4}n^2+O(n)}$, the number of subgroups of \mathbb{Z}_p^n . Our Theorem 11 then provides the following improvement, whose proof is in Section 12.

Theorem 16. *Let P be as above. Then the number of maximal abelian subgroups of P is upper bounded by $p^{\frac{1}{6}n^2+O(n)}$.*

Recall that the proof of Theorem 11 starts by following the strategy of Wood’s proof [Woo11] of bounding the number of maximal independent sets on a graph. We view this as an interesting and somewhat unexpected example of transferring techniques from graph theory to group theory.

We also initiate a quantum variant of the theory in Section 13. There, the objects are a special type of quantum channels, and isotropic spaces and isotropic decompositions are defined on the Kraus operators of such channels. Furthermore, we require the isotropic decompositions to be orthogonal. One can then transform classical connected graphs into such channels, and prove an analogue of Theorem 3. More surprisingly, we also obtain an efficient isotropic 2-decomposition algorithm, as an analogue of Theorem 6, by resorting to the recent development on the periodicity of quantum Markov chains [GFY18].

We then present an information theoretic interpretation for isotropic spaces in the context of quantum error correction. Briefly speaking, from the viewpoint of certain natural generalizations of quantum gate fidelities [NC02], isotropic spaces can be viewed as the opposite structure of noiseless subspaces (Proposition 48), which have been studied intensively in quantum error correction [KL97, Lid14]. Indeed, noiseless subspaces are shelters for the information residing in them under quantum noise, while the information in an isotropic space would be completely destroyed by quantum noise.

Let us conclude this subsection with a remark on these applications. After building a bridge, we expect it to serve as a two-way street between the two sides. However, in reality there is usually more traffic in one direction than the other. For example, the traffic between perfect matchings and full-rank matrices mostly goes from the algebra side to the combinatorial side, e.g., the randomized NC algorithm for perfect matchings [Lov79]. The traffic between shrunk subsets and shrunk subspaces mostly goes in the other direction, e.g., linear algebraic analogues of augmenting paths [IKQS15] and scaling [Gur04]. So far, our applications in this work mostly go in the direction from combinatorics to algebra, following the pattern of the shrunk subset vs. shrunk subspace case. It will be very interesting to explore implications in the other direction in the future.

2.5 Outlook

Summary of our contributions. The concepts of isotropic spaces and isotropic decompositions for alternating bilinear maps are classical, with natural interpretations in group theory and manifold theory. Our key new insight is that they can be viewed and studied as linear algebraic analogues of independent sets and vertex colorings. This insight leads us to study algorithmic and mathematical problems about isotropic spaces and isotropic decompositions, by drawing inspirations from results and techniques from graph theory and algorithm study. The techniques used to address the problems range from combinatorics, to algebra, and to quantum information.

We believe that this investigation is fruitful, for the following reasons.

1. First, it discloses new algorithmic and mathematical questions. For example, in Section 2.2 we proposed and studied upper bounding the number of maximal isotropic spaces, and exact exponential-time algorithms for isotropic decomposition numbers over finite fields.
2. Second, the results obtained have concrete applications to other mathematical and computational disciplines. For example, in Section 2.4, we described the applications of our results to finite groups and quantum information.
3. Third, it sheds new lights on known results from different research directions. For example, in Section 2.1 we compared Turán’s extremal graph result with Buhler, Gupta, and Harris’ algebraic geometric result.

This investigation then lays the foundation for yet another bridge between graphs and alternating matrix spaces, adding to the classical ones established by Tutte and Lovász.

Open ends. Several interesting open problems have been mentioned before, and here we give a summary and propose some new ones.

1. By Theorem 11, the number of maximal isotropic spaces of $\mathcal{A} \leq \Lambda(n, q)$ is upper bounded by $f(n, q) = q^{\frac{1}{6}n^2 + O(n)}$. There exists an alternating matrix space with $g(n, q) = q^{\frac{1}{8}n^2 + \Omega(n)}$ many maximal isotropic spaces (see Section 2.2). Either improve the current upper bound $f(n, q)$, or construct an alternating matrix space with more than $g(n, q)$ maximal isotropic spaces. Note that resolving this problem would lead to a sharp bound on the number of maximal abelian subgroups of p -groups of class 2 and exponent p .
2. Improve the exact exponential-time algorithm for computing the isotropic decomposition number for $\mathcal{A} \leq \Lambda(n, q)$ in Theorem 12. An interesting question is whether the strategy in [BHK09] can be adapted here. The results in [BHK⁺16] should be useful in this context.
3. Despite Theorem 10, the complexities of deciding whether an alternating matrix space has an isotropic space of dimension ≥ 2 are not clear over various fields. Even over \mathbb{Q} , our proof for Theorem 10 relies on a special case of the underlying existential singularity problem for matrix spaces (see Section 8.1), so it is left open even for the general case of that problem over \mathbb{Q} .
4. Investigate the behaviours of the isotropic and isotropic decomposition numbers in the linear algebraic Erdős-Rényi model [LQ17, Bol01].
5. Improve the dependence of the isotropic decomposition number on the maximum degree, or the dimension of the alternating matrix space (see Proposition 14). Note that this problem has motivations from classical geometry (see the discussions before Proposition 14).

3 Preliminaries

Notation. For $n \in \mathbb{N}$, $[n] := \{1, \dots, n\}$. We use \uplus for disjoint union of sets. The base of the logarithm is 2 unless otherwise stated.

Let \mathbb{F} be a field. We use \mathbb{F}^n to denote the vector space of *column* vectors of length n over \mathbb{F} . The standard basis of \mathbb{F}^n consists of vectors e_1, \dots, e_n , where e_i is the vector with the i th entry being 1, and other entries being 0. The linear span of several vectors or matrices is denoted by $\langle \cdot \rangle$.

For $n, d \in \mathbb{N}$, let $M(n \times d, \mathbb{F})$ be the linear space of $n \times d$ matrices over \mathbb{F} , and $GL(n \times d, \mathbb{F})$ the set of $n \times d$ matrices over \mathbb{F} of rank $\min(n, d)$. We also let $M(n, \mathbb{F}) := M(n \times n, \mathbb{F})$, and $GL(n, \mathbb{F}) := GL(n \times n, \mathbb{F})$. Dimension- d subspaces of \mathbb{F}^n will be understood as represented by elements from $GL(n \times d, \mathbb{F})$. Given $A \in M(n \times d, \mathbb{F})$, the transpose of A is denoted by $A^t \in M(d \times n, \mathbb{F})$. For convenience, we sometimes write a vector v in \mathbb{F}^n as $v = (v_1, \dots, v_n)^t$.

Depending on the context, $\mathbf{0}$ may denote either the zero space, a zero vector, or a zero matrix. The identity matrix in $M(n, \mathbb{F})$ is denoted by I_n ; we may drop the subscript n when it is understood from the context. Given a matrix $A \in M(n, \mathbb{F})$, its kernel and image are denoted by $\ker(A)$ and $\text{im}(A)$, respectively. For $U \leq \mathbb{F}^n$, the image of U under A is denoted by $A(U)$.

Linear algebra. Given $U \leq \mathbb{F}^n$, a complementary subspace, or just a complement, is some $V \leq \mathbb{F}^n$ such that $V \cap U = \mathbf{0}$, and $\langle U \cup V \rangle = \mathbb{F}^n$. Note that complement subspaces of U are not unique. Indeed, the number of complements of a dimension- d subspace $U \leq \mathbb{F}_q^n$ is $q^{d(n-d)}$. The space orthogonal to U is $\{v \in \mathbb{F}^n : \forall u \in U, v^t u = 0\}$. (Over \mathbb{C} , the conjugate transpose is used.) Note that the space orthogonal to U is not necessarily a complement to U .

Some algebra. The following class of finite groups is of main interest in this article. Let p be a prime > 2 . A group G belongs to the class of *p -groups of class 2 and exponent p* , if G satisfies the following conditions: (1) G is of prime power order, (2) the commutator subgroup $[G, G]$ is contained in the centre $Z(G)$, and (3) every group element $g \in G$ satisfying $g^p = 1$.

Let \mathbb{K}/\mathbb{F} be a quadratic field extension. That is, \mathbb{K} is of dimension 2 as a \mathbb{F} -vector, and every element in \mathbb{K} can be represented as $a + b\gamma$ for some $\gamma \in \mathbb{K} \setminus \mathbb{F}$. The map sending $a + b\gamma$ to $a - b\gamma$ is a field automorphism, called the quadratic involution for this field extension. A basic example is when $\mathbb{K} = \mathbb{C}$, $\mathbb{F} = \mathbb{R}$, and $\gamma = i$, the imaginary unit. For a matrix $A \in M(n \times d, \mathbb{K})$, A^\dagger denotes the conjugate transpose of A .

On matrix spaces. Given a matrix space $\mathcal{A} \leq M(s \times t, \mathbb{F})$, the image of $U \leq \mathbb{F}^t$ under \mathcal{A} is $\mathcal{A}(U) := \langle \cup_{A \in \mathcal{A}} A(U) \rangle$. The dimension of \mathcal{A} is denoted by $\dim(\mathcal{A})$. The (maximum) rank of \mathcal{A} is $\text{rk}(\mathcal{A}) := \max\{\text{rk}(A) : A \in \mathcal{A}\}$. Let $\mathcal{B} \leq M(s \times t, \mathbb{F})$ be another matrix space. We say that \mathcal{A} and \mathcal{B} are *equivalent*, if there exist $C \in GL(s, \mathbb{F})$ and $D \in GL(t, \mathbb{F})$, such that $\mathcal{A} = CBD := \{CBD : B \in \mathcal{B}\}$. When working with \mathcal{A} , an equivalence transformation is meant to left multiply \mathcal{A} with some $C \in GL(s, \mathbb{F})$ and right multiply it with some $D \in GL(t, \mathbb{F})$.

On alternating matrices. Let $A, B \in \Lambda(n, \mathbb{F})$. We say that A and B are *isometric*, if there exists $T \in GL(n, \mathbb{F})$, such that $A = T^t B T$. Given a dimension- d $U \leq \mathbb{F}^n$ represented by $T \in GL(n \times d, \mathbb{F})$, the *restriction* of A to U by T , denoted as $A|_{U, T}$, is $T^t A T \in \Lambda(d, \mathbb{F})$. The radical of A is the subspace $\{u \in \mathbb{F}^n : \forall v \in \mathbb{F}^n, v^t A u = 0\}$, which is just $\ker(A)$. The rank of A is always even. If

$\text{rk}(A) = 2r$, then A is isometric to $\begin{bmatrix} \mathbf{0} & I_r & \mathbf{0} \\ -I_r & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix}$ (see e.g. [Lan02, Chap. XV, Sec. 8]). We say

that $A \in \Lambda(n, \mathbb{F})$ is *non-degenerate*, if A is full-rank (so n is even).

On alternating matrix spaces. Let $\mathcal{A}, \mathcal{B} \leq \Lambda(n, \mathbb{F})$. We say that \mathcal{A} and \mathcal{B} are *isometric*, if there exists $T \in \text{GL}(n, \mathbb{F})$, such that $\mathcal{A} = T^t \mathcal{B} T := \{T^t B T : B \in \mathcal{B}\}$. Given a dimension- d $U \leq \mathbb{F}^n$ represented by $T \in \text{GL}(n \times d, \mathbb{F})$, the *restriction* of \mathcal{A} on U via T is $\mathcal{A}|_{U, T} := \{T^t A T : A \in \mathcal{A}\} \leq \Lambda(d, \mathbb{F})$. When it does not cause confusion, we may not write T explicitly, and just say the restriction of \mathcal{A} to U , denoted by $\mathcal{A}|_U$. This corresponds to the concept of induced subgraphs in graph theory. Indeed, we see that U is an isotropic space if and only if $\mathcal{A}|_U$ is the zero (alternating matrix) space. Given $v \in \mathbb{F}^n$, the *radical of v in \mathcal{A}* , denoted as $\text{rad}_{\mathcal{A}}(v)$, is $\{u \in \mathbb{F}^n : \forall A \in \mathcal{A}, u^t A v = 0\}$ which is a subspace of \mathbb{F}^n . Elements in $\text{rad}_{\mathcal{A}}(v)$ correspond to non-neighbours in graph theory. The *codegree* of v in \mathcal{A} , denoted as $\text{codeg}_{\mathcal{A}}(v)$, is $\dim(\text{rad}_{\mathcal{A}}(v))$. Note that $\text{codeg}_{\mathcal{A}}(v) \geq 1$ for non-zero v , as $v \in \text{rad}_{\mathcal{A}}(v)$. The *degree* of v in \mathcal{A} is $\text{deg}_{\mathcal{A}}(v) := n - \text{codeg}_{\mathcal{A}}(v)$. More generally, for $U \leq \mathbb{F}^n$, $\text{rad}_{\mathcal{A}}(U) = \{v \in \mathbb{F}^n : \forall u \in U, \forall A \in \mathcal{A}, u^t A v = 0\}$. When \mathcal{A} is clear from the context, we may drop the subscript \mathcal{A} in $\text{rad}_{\mathcal{A}}(v)$, $\text{codeg}_{\mathcal{A}}(v)$, $\text{deg}_{\mathcal{A}}(v)$, etc.. It is easy to see that for any $v \in \text{rad}(U)$, we have $U \leq \text{rad}(v)$, or in other words, $U \leq \text{rad}(\text{rad}(U))$.

A vector $v \in \mathbb{F}^n$ is called *isolated* in \mathcal{A} , if for any $A \in \mathcal{A}$, $Av = \mathbf{0}$, which is equivalent to say that $\text{deg}(v) = 0$. This corresponds to the concept of isolated vertices in graph theory. The *radical* of \mathcal{A} , $\text{rad}(\mathcal{A})$, is the subspace of \mathbb{F}^n consisting of all isolated vectors. We say that \mathcal{A} is *non-degenerate*, if $\text{rad}(\mathcal{A}) = \mathbf{0}$, and *degenerate* otherwise. If $\mathcal{A} \leq \Lambda(n, \mathbb{F})$ is degenerate with $\dim(\text{rad}(\mathcal{A})) = d > 0$, then \mathcal{A} is isometric to \mathcal{A}' where each $A \in \mathcal{A}'$ is of the form $\begin{bmatrix} A' & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$, where $A' \in \Lambda(n - d, \mathbb{F})$.

Sets, tuples, and spaces. Let $\mathcal{A} \leq \Lambda(n, \mathbb{F})$ be given by a linear basis $A_1, \dots, A_m \in \Lambda(n, \mathbb{F})$. We can collect them as a set $\mathbf{A} = \{A_1, \dots, A_m\} \subseteq \Lambda(n, \mathbb{F})$. Sometimes it is also useful to impose an order on them, and form a tuple $\mathbf{A} = (A_1, \dots, A_m) \in \Lambda(n, \mathbb{F})^m$. We shall use calligraphic fonts for spaces, bold fonts for tuples, and sans serif fonts for sets.

Suppose $\mathcal{A} \leq \Lambda(n, \mathbb{F})$ is given by a linear basis $A_1, \dots, A_m \in \Lambda(n, \mathbb{F})$. Then it is clear that, given $U \leq \mathbb{F}^n$, for any $u, u' \in U$ and any $A \in \mathcal{A}$, $u^t A u' = 0$, if and only if for any $u, u' \in U$ and any $i \in [m]$, $u^t A_i u' = 0$. We therefore can define isotropic spaces, and isotropic decompositions, for sets or tuples of alternating matrices. In particular, since alternating matrix tuples represent alternating bilinear maps naturally (see Remark 17), this observation suggests that isotropic spaces and decompositions for alternating matrix spaces and for alternating bilinear maps are basically the same object. Furthermore, many, though not all, concepts introduced in Section 3 about alternating matrix spaces can be translated naturally to alternating bilinear maps, including degrees, degeneracy, radicals, etc.. (Indeed, some notions there are actually borrowed from alternating bilinear maps.) On the other hand, note that the maximum rank in \mathcal{A} is more naturally associated with the space perspective. More discussions on the relation between these two notions are in Remark 17.

Therefore, in the context of isotropic spaces and decompositions, the choices between alternating matrix spaces and tuples usually do not matter much. The tuple perspective is more natural from the algorithm perspective, because the input of an alternating matrix space to an algorithm is usually an ordered basis. The space perspective is more natural for forming the analogy with graphs, and more naturally allows for some more notions including the maximum rank, which is important in e.g. the proof of Theorem 10. The set perspective will be used in the quantum variant of the theory in Section 13. Therefore, it is best to keep all three perspectives in mind, and see how they fit into our problems.

Remark 17. It is easy to realize that alternating matrix spaces are closely related to alternating bilinear maps.

For this, we first recall the relation between alternating bilinear maps and alternating matrix tuples. Let $\phi : U \times U \rightarrow V$ be an alternating bilinear map, that is, for any $u \in U$, $\phi(u, u) = \mathbf{0}$. Fix bases of U and V , so that $U \cong \mathbb{F}^n$ and $V \cong \mathbb{F}^m$. Then ϕ can be represented by an m -tuple of alternating matrices $(A_1, \dots, A_m) \in \Lambda(n, \mathbb{F})^m$, such that $\phi(u, u') = (u^t A_1 u', \dots, u^t A_m u')^t$. Conversely, given an m -tuple of alternating matrices, one can define an alternating bilinear map as such. Two alternating bilinear maps $\phi, \psi : U \times U \rightarrow V$ are isometric, if there exist $A \in \text{GL}(U)$, $B \in \text{GL}(V)$, such that $\phi = B \circ \psi \circ A$. (Some authors prefer to call this isometric as pseudo-isometric [BW12].)

Let $\mathcal{A} \leq \Lambda(n, \mathbb{F})$. Let $\mathbf{A} = (A_1, \dots, A_m) \in \Lambda(n, \mathbb{F})^m$ be an ordered basis of \mathcal{A} . Then \mathbf{A} defines an alternating bilinear map $\phi_{\mathbf{A}} : \mathbb{F}^n \times \mathbb{F}^n \rightarrow \mathbb{F}^m$ as above. While different choices of ordered bases give different alternating bilinear maps, it is easy to see that ordered bases from isometric alternating matrix spaces give isometric alternating bilinear maps.

3.1 Computational models

We will work with two computational models, depending on the problems. The first model may be called the exact model; see e.g. [Lov86]. This is the model to work with, if field extensions are unavoidable. In this model, input matrices or vectors are over a field \mathbb{E} where \mathbb{E} is a finite field extension over its prime field \mathbb{F} . Recall that a prime field is a field containing no proper subfields, so \mathbb{F} is either a field of prime order or \mathbb{Q} . Suppose $\dim_{\mathbb{F}}(\mathbb{E}) = d$. Then \mathbb{E} is an extension of \mathbb{F} by a single generating element α . We represent α by the minimal polynomial of α over \mathbb{F} , and an isolating interval for α in the case of \mathbb{R} , or an isolating rectangle for α in the case of \mathbb{C} . Note that from this representation, one can approximate the numerical value of α arbitrarily closely. When we say that we work over \mathbb{R} or \mathbb{C} , the input is given as over some algebraic number field \mathbb{E} in \mathbb{R} or \mathbb{C} ; recall that an algebraic number field is a finite degree extension of \mathbb{Q} . The algorithm is allowed to work with extension fields of \mathbb{E} in \mathbb{R} or \mathbb{C} , as long as the extension degrees are polynomially bounded.

The second model may be called the arithmetic model. In this model, only basic field operations are performed, and the issue of working with different field extensions does not arise. Still, over number fields we will be concerned with the bit complexities, though it is possible that we may be able to only bound the number of arithmetic steps, but not the bit complexities.

We shall mostly work with \mathbb{F}_q , \mathbb{Q} , \mathbb{R} , and \mathbb{C} in this article, though some results extend to number fields naturally. Sometimes, we make further restrictions like requiring q to be odd, or the input to be integral.

4 Basic facts and properties

In this section we collect some basic results about isotropic spaces and isotropic decompositions.

On the definitions. The following is a somewhat more intuitive definition of these two notions.

Let $\mathcal{A} \leq \Lambda(n, \mathbb{F})$, and let $U \leq \mathbb{F}^n$ be an isotropic space. Suppose $d = \dim(U)$. Then form a change of basis matrix $T \in \text{GL}(n, \mathbb{F})$, such that its first d columns form a basis of U , and the rest columns together with the first d ones span \mathbb{F}^n . Then for any $A \in \mathcal{A}$, we see that $T^t A T$ is of the form $\begin{bmatrix} \mathbf{0} & B \\ -B^t & C \end{bmatrix}$ where $\mathbf{0}$ is of size $d \times d$. It is not hard to see that \mathcal{A} has a dimension- d isotropic space if and only if there exists such a change of basis matrix T .

Similarly, \mathcal{A} has an isotropic c -decomposition, if and only if there exist $T \in \text{GL}(n, \mathbb{F})$, and $d_1, \dots, d_c \in \mathbb{Z}^+$ with $\sum_i d_i = n$, such that for any $A \in \mathcal{A}$, $T^t A T =$

$$\begin{bmatrix} \mathbf{0} & A_{1,2} & \dots & A_{1,c} \\ -A_{1,2}^t & \mathbf{0} & \dots & A_{2,c} \\ \vdots & \vdots & \ddots & \vdots \\ -A_{1,c}^t & -A_{2,c}^t & \dots & \mathbf{0} \end{bmatrix},$$

where the i th $\mathbf{0}$ on the diagonal is of size d_i .

Computing the radical. For many problems about isotropic spaces, given a degenerate $\mathcal{A} \leq \Lambda(n, \mathbb{F})$, usually it is possible to reduce to the non-degenerate case. This is facilitated by the fact that the radical of \mathcal{A} is easy to compute.

Observation 18. *Suppose $\mathcal{A} = \langle A_1, \dots, A_m \rangle \leq \Lambda(n, \mathbb{F})$ is given by a linear basis. There is a polynomial-time algorithm that computes a linear basis of $\text{rad}(\mathcal{A})$.*

Proof. Observe that $\text{rad}(\mathcal{A}) = \bigcap_{i=1}^m \ker(A_i)$. Computing $\ker(A_i)$ and the intersection of $\ker(A_i)$'s are standard linear algebraic tasks that can be performed as stated. \square

Isotropic spaces and radicals of subspaces. Let $\mathcal{A} \leq \Lambda(n, \mathbb{F})$. Recall that for $U \leq \mathbb{F}^n$, we defined $\text{rad}_{\mathcal{A}}(U) = \{v \in \mathbb{F}^n : \forall u \in U, v^t A u = 0\}$. The following observation is immediate.

Observation 19. *Let $U \leq \mathbb{F}^n$ and $\mathcal{A} \leq \Lambda(n, \mathbb{F})$. Then we have the following.*

1. U is an isotropic space of \mathcal{A} if and only if $U \subseteq \text{rad}(U)$.
2. U is a maximal isotropic space of \mathcal{A} if and only if $U = \text{rad}(U)$.

An easy application of Observation 19 gives a greedy algorithm for computing one maximal isotropic space.

Proposition 20. *Given a matrix space $\mathcal{A} \leq \Lambda(n, \mathbb{F})$, a maximal isotropic space can be computed in polynomially many arithmetic steps.*

Proof. We first present the algorithm. Recall that e_i is the i th standard basis vector of \mathbb{F}^n .

1. Let $U = \langle e_1 \rangle$.
2. While $U \subsetneq \text{rad}(U)$:
 - (a) Take any $u \in \text{rad}(U) \setminus U$.
 - (b) $U \leftarrow \langle U, u \rangle$.
3. Output U .

To see the correctness, note that in Step (2.a), by the choice of u , we have $U \subseteq \text{rad}(\langle U, u \rangle)$ and $u \in \text{rad}(\langle U, u \rangle)$, so $\langle U, u \rangle$ is an isotropic space by Observation 19 (1). In Step (3), U satisfies $U = \text{rad}(U)$, so U is maximal by Observation 19 (2).

To see the running time, note that the while loop will be executed by at most n times, since $\dim(U)$ increases by 1 in each execution. Each step involves basic linear algebraic computations which require only polynomially many arithmetic operations. \square

Over \mathbb{R} or \mathbb{C} , the bit sizes in the above algorithm may blow up, at least with a straightforward implementation, due to the iterative computations of the radicals.

On field extensions. Let \mathbb{K} be an extension field of \mathbb{F} . Given $\mathcal{A} \leq \Lambda(n, \mathbb{F})$, we let $\mathcal{A}_{\mathbb{K}} \leq \Lambda(n, \mathbb{K})$ be the alternating matrix spaces when we allow linear combinations over \mathbb{K} , or in other words, $\mathcal{A}_{\mathbb{K}} = \mathcal{A} \otimes_{\mathbb{F}} \mathbb{K}$. We may write \mathcal{A} as $\mathcal{A}_{\mathbb{F}}$ for further distinction. For $\mathcal{A}_{\mathbb{K}}$ we allow for isotropic spaces to come from \mathbb{K}^n . Since isotropic spaces and c -decompositions of $\mathcal{A}_{\mathbb{F}}$ naturally give isotropic spaces and c -decompositions of $\mathcal{A}_{\mathbb{K}}$, we have the following (see also [Gel17, Lemma 6]).

Proposition 21. *Let $\mathcal{A}_{\mathbb{F}} \leq \Lambda(n, \mathbb{F})$ and $\mathcal{A}_{\mathbb{K}} \leq \Lambda(n, \mathbb{K})$ be as above. Then $\alpha(\mathcal{A}_{\mathbb{F}}) \leq \alpha(\mathcal{A}_{\mathbb{K}})$, and $\chi(\mathcal{A}_{\mathbb{F}}) \geq \chi(\mathcal{A}_{\mathbb{K}})$.*

In Proposition 21, the inequalities for $\alpha(\cdot)$ and $\chi(\cdot)$ could be strict, as shown by Buhler, Gupta, and Harris [BGH87, pp. 277]. Recall that in Section 2.1 we defined $\alpha(\mathbb{F}, n, m) = \{\alpha(\mathcal{A}) : \mathcal{A} \leq \Lambda(n, \mathbb{F}), \dim(\mathcal{A}) = m\}$, and Buhler et al. showed that $\alpha(\mathbb{F}, n, m) \leq \lceil \frac{m+2n}{m+2} \rceil$, when $m > 1$ and the characteristic of \mathbb{F} is not 2. Furthermore the equality can be attained for algebraically closed fields. Buhler et al. then demonstrated examples over \mathbb{F}_q and \mathbb{Q} , for which the inequality is strict. Consider the example over \mathbb{Q} , which is some $\mathcal{A} \leq \Lambda(n, \mathbb{Q})$ of dimension n . They show that $\alpha(\mathcal{A}_{\mathbb{Q}}) = 1$ and $\alpha(\mathcal{A}_{\mathbb{C}}) = 2$, which is equivalent to that $\chi(\mathcal{A}_{\mathbb{Q}}) = n$ and $\chi(\mathcal{A}_{\mathbb{C}}) \leq n - 1$. This gives the desired separations. Some interesting discussions on \mathbb{R} vs \mathbb{C} can be found in [BGH87, Sec. 3].

4.1 Isotropic spaces and decompositions in group theory and manifold theory

In this subsection, we explain the origins of isotropic spaces and decompositions for alternating matrix spaces in group theory and manifold theory. These are classical, so the purpose here is to provide references for interested readers who have not met with these before.

Group theory. Let P be a p -group of class 2 and exponent p of order p^ℓ . We have that the commutator subgroup $[P, P] \cong \mathbb{Z}_p^m$, and the commutator quotient $P/[P, P] \cong \mathbb{Z}_p^n$. The commutator bracket then gives an alternating bilinear map $\phi : P/[P, P] \times P/[P, P] \rightarrow [P, P]$, or, after fixing bases of $P/[P, P]$ and $P/[P, P]$, an alternating bilinear map $\phi : \mathbb{Z}_p^n \times \mathbb{Z}_p^n \rightarrow \mathbb{Z}_p^m$. On the other hand, given an alternating bilinear map $\phi : \mathbb{F}_p^n \times \mathbb{F}_p^n \rightarrow \mathbb{F}_p^m$, one can construct a p -group of class 2 and exponent p , P_ϕ , called the *Baer group* corresponding to ϕ , as follows. The group elements are $(v, u) \in \mathbb{F}_p^n \times \mathbb{F}_p^m$, and the group multiplication \circ is by $(v_1, u_1) \circ (v_2, u_2) = (v_1 + v_2, u_1 + u_2 + \frac{1}{2} \cdot \phi(v_1, v_2))$. This sets up a two-way correspondence between p -groups of class 2 and exponent p and alternating bilinear maps.

Having set up the connection, let us see the group-theoretic interpretations of isotropic spaces and decompositions. The following is classical: an isotropic space of ϕ corresponds to a normal abelian subgroup of P containing the commutator subgroup (see e.g. [Alp65b]). More recently, Lewis and Wilson proposed the concept of a hyperbolic pair of P [LW12], which just consists of two normal abelian subgroups of P which together generate P , and whose intersection equals $[P, P]$. A natural generalization is then the following. A hyperbolic c -system of P consists of c normal abelian subgroups A_1, A_2, \dots, A_c , such that P is generated by A_i , and for any $i, j \in [c]$, $i \neq j$, $A_i \cap A_j = [P, P]$. A hyperbolic c -system then naturally corresponds to an isotropic c -decomposition of ϕ .

Manifold theory. We then turn to the manifold theory side. We shall just walk through some examples of the connection, and point the interested reader to the survey [Dim08] for more detailed information.

Let M be a compact Kähler manifold. Let $H^i(M; \mathbb{C})$ be the i th cohomology group of M with coefficients in \mathbb{C} . The cup product $\smile: H^1(M; \mathbb{C}) \times H^1(M; \mathbb{C}) \rightarrow H^2(M; \mathbb{C})$ is a skew-symmetric bilinear map. Then as an application of results by Castelnuovo and de Franchis, Catanese [Cat91] showed that there exists a constant holomorphic map $f: M \rightarrow C$, where C is a curve of genus $g \geq 2$, if and only if, \smile has a dimension- g maximal isotropic space. Catanese went on to generalize this connection much further in [Cat91].

Let M be a smooth closed orientable n -dimensional manifold. Let $H^i(M; \mathbb{Q})$ be the i th cohomology group of M with coefficients in \mathbb{Q} . Gelbukh showed that an isotropic space of \smile corresponds to a geometric structure on M [Gel17, Lemma 10, Definition 11, Theorem 13], called an isotropic system, which consists of smooth closed orientable connected codimension-one submanifolds that are homologically non-intersecting, homologically independent, and intersecting transversely. The isotropic index defined in [Gel17] for an alternating bilinear map is just the isotropic number introduced in Definition 2. In particular, our Theorem 3 shows that computing this isotropic index is NP-hard. The relations of isotropic indices with other basic notions in manifold cohomology, including the first Betti number and the co-rank of the fundamental group, are also studied there.

5 Proof of Theorem 3

Recall that we have a graph $G = ([n], E)$, and an alternating matrix space \mathcal{A}_G which is spanned by those elementary alternating matrices $A_{i,j}$ where $\{i, j\} \in E$.

(1) Isotropic spaces and independent sets. We need to show that $G = ([n], E)$ has a size- s independent set if and only if \mathcal{A}_G has a dimension- s isotropic space.

For the only if direction, let $T = \{i_1, \dots, i_s\}$ be a size- s independent set of G . Let U be the subspace of \mathbb{F}^n spanned by e_{i_1}, \dots, e_{i_s} ; recall that e_i denotes the i th standard basis vector of \mathbb{F}^n . It is easy to verify that U is an isotropic space of \mathcal{A} of dimension s .

For the if direction, let $U = \langle u_1, \dots, u_s \rangle$ be a dimension- s isotropic space of \mathcal{A}_G . We form an $n \times s$ matrix U such that the i th column of U is u_i , that is, $U = [u_1, \dots, u_s] \in \text{GL}(n \times s, \mathbb{F})$. Suppose $U^t = [w_1, \dots, w_n] \in \text{GL}(s \times n, \mathbb{F})$, $w_i \in \mathbb{F}^s$. Since U is of rank s , there exist integers i_1, \dots, i_s , $1 \leq i_1 < \dots < i_s \leq n$, such that w_{i_1}, \dots, w_{i_s} are linearly independent. We now claim that $\{i_1, \dots, i_s\}$ forms an independent set of the original graph $G = ([n], E)$. If not, suppose $\{i_j, i_{j'}\}$, $1 \leq j < j' \leq s$, is in E . As U is an isotropic space of \mathcal{A}_G , we have that for any $\{k, \ell\} \in E$, $U^t A_{k,\ell} U = \mathbf{0}$. As $A_{k,\ell} = e_k e_\ell^t - e_\ell e_k^t$, we have

$$U^t A_{k,\ell} U = U^t (e_k e_\ell^t - e_\ell e_k^t) U = w_k w_\ell^t - w_\ell w_k^t = 0,$$

which implies that w_k and w_ℓ are linearly dependent. It follows that w_{i_j} and $w_{i_{j'}}$ are linearly dependent. We then arrive at a contradiction.

(2) Isotropic decompositions and vertex colorings. We need to show that $G = ([n], E)$ has a vertex c -coloring if and only if \mathcal{A}_G has an isotropic c -decomposition.

For the only if direction, assume G has a vertex c -coloring, and let $[n] = T_1 \uplus T_2 \uplus \dots \uplus T_c$ be a partition of $[n]$ into disjoint union of independent sets. Suppose $|T_j| = t_j$, and $T_j = \{i_{j,1}, \dots, i_{j,t_j}\} \subseteq [n]$. Let $U_j = \langle e_{i_{j,1}}, \dots, e_{i_{j,t_j}} \rangle \leq \mathbb{F}^n$, so $\mathbb{F}^n = U_1 \oplus U_2 \oplus \dots \oplus U_c$. By (1), every U_j is an isotropic space. This gives an isotropic c -decomposition of \mathcal{A}_G .

For the if direction, let $\mathbb{F}^n = U_1 \oplus U_2 \oplus \cdots \oplus U_c$ be an isotropic c -decomposition. Let $d_i = \dim(U_i)$, and $b_i = \sum_{j=1}^i d_j$, for $i \in [c]$. Set $b_0 = 0$. Let $P = [p_1, \dots, p_n]$ be an $n \times n$ invertible matrix, where $p_i \in \mathbb{F}^n$, such that $p_{b_{i-1}+1}, \dots, p_{b_i}$ form a basis of U_i . By abuse of notation, we also let $U_i = [p_{b_{i-1}+1}, \dots, p_{b_i}] \in \text{GL}(n \times d_i, \mathbb{F})$, so $P = [U_1, \dots, U_c]$. Let $W_i = U_i^t = [w_{i,1}, \dots, w_{i,n}] \in \text{GL}(d_i \times n, \mathbb{F})$. Since U_i is an isotropic space, by (1), we know that for any $\{k, \ell\} \in E$, $w_{i,k}$ and $w_{i,\ell}$ are linearly dependent.

We then use the following simple linear algebraic result, which is a consequence of the Laplacian expansion. For $A, B \subseteq [n]$ of the same size, we let $C|_{A,B}$ to denote the submatrix of C with row indices from A and column indices from B .

Lemma 22. *Let $P = [U_1, \dots, U_c] \in \text{GL}(n, \mathbb{F})$, where $U_i \in \text{GL}(n \times d_i, \mathbb{F})$. Then there exists a partition of $[n] = T_1 \uplus T_2 \uplus \cdots \uplus T_c$, where $|T_i| = d_i$, such that $\forall i \in [c]$, $\text{rk}(U_i|_{T_i, [d_i]}) = d_i$.*

We claim that the partition of $[n] = T_1 \uplus T_2 \uplus \cdots \uplus T_c$ from Lemma 22 gives a vertex c -coloring of G . To see this, observe that, the condition $\text{rk}(U_i|_{T_i, [d_i]}) = d_i$ is equivalent to that the vectors $w_{i,j}$, $j \in T_i$, are linearly independent. This implies that G cannot have edges of the form $\{k, \ell\}$ where $k, \ell \in T_i$, as otherwise $w_{i,k}$ and $w_{i,\ell}$ would be linearly dependent. Hence T_i is an independent set for any $i \in [c]$. This completes the proof of the second part of Theorem 3.

6 An exposition of the proof of Theorem 6

As mentioned in Section 2.1, we give an exposition of the proof of Theorem 6 for \mathbb{F}_q in [BMW17], using some ingredients from [IQ19] to handle \mathbb{R} and \mathbb{C} . The main purpose is to give the reader a flavor of how the so-called $*$ -algebra technique is applied in this context. We could not give all the details here, as that would be too long and unnecessary; the interested reader may wish to go to [Wil09a, BW12, IQ19], which contain detailed proofs for using $*$ -algebras to solve several closely related problems.

Recall that we are given $\mathcal{A} = \langle A_1, \dots, A_m \rangle \leq \Lambda(n, \mathbb{F})$, and our goal is to find a non-trivial direct sum decomposition $\mathbb{F}^n = U_1 \oplus U_2$, such that $\mathcal{A}|_{U_i} = \mathbf{0}$ for $i = 1, 2$. We first reduce to the non-degenerate setting as follows. Suppose \mathcal{A} is degenerate. Then let T be a complement subspace of the radical, and set $\mathcal{A}' = \mathcal{A}|_T$. Note that \mathcal{A}' is non-degenerate. Then it is not hard to verify that \mathcal{A} admits an isotropic 2-decomposition if and only if \mathcal{A}' admits an isotropic 2-decomposition.

In the following we assume that \mathcal{A} is non-degenerate. Let $\mathbf{A} = (A_1, \dots, A_m) \in \Lambda(n, \mathbb{F})^m$. The adjoint algebra of \mathbf{A} is defined as $\text{Adj}(\mathbf{A}) := \{D \in \text{M}(n, \mathbb{F}) : \exists B \in \text{M}(n, \mathbb{F}), \forall i \in [m], B^t A_i = A_i D\}$. Since \mathbf{A} is non-degenerate, if such B exists, then it is unique. Then a natural involution (an anti-automorphism of order at most 2) on $\text{Adj}(\mathbf{A})$ is to send $D \in \text{Adj}(\mathbf{A})$ to this unique B satisfying $B^t A_i = A_i D$ for any $i \in [m]$, also denoted as D^* . The adjoint algebra is the key device for the algorithm.

We now translate the isotropic 2-decomposition problem for \mathcal{A} , and therefore \mathbf{A} , to a problem about $\text{Adj}(\mathbf{A})$. Any direct sum decomposition $\mathbb{F}^n = U_1 \oplus U_2$ can be encoded as a projection matrix $P \in \text{M}(n, \mathbb{F})$, that is, $P^2 = P$, $\text{im}(P) = U_1$, $\text{ker}(P) = U_2$. The key observation in [BMW17] is that, P corresponds to an isotropic 2-decomposition if and only if $P \in \text{Adj}(\mathbf{A})$ and $P^* = I - P$. This means that we need to search for an idempotent P in $\text{Adj}(\mathbf{A})$ satisfying $P^* = I - P$. Following [BMW17], we call such an idempotent a hyperbolic idempotent.

To do that, we utilize the $*$ -algebra structure of $\text{Adj}(\mathbf{A})$. For this, we recall in a nutshell the structure of $*$ -algebras. Let \mathfrak{A} be a $*$ -algebra. The Jacobson radical of \mathfrak{A} , denoted by $\text{rad}(\mathfrak{A})$,

is the largest nilpotent ideal of \mathfrak{A} , and it is invariant under $*$. The factor algebra $\mathfrak{A}/\text{rad}(\mathfrak{A})$ is semi-simple, namely it is a direct sum of simple algebras. Let $\mathfrak{A}/\text{rad}(\mathfrak{A}) \cong S_1 \oplus \cdots \oplus S_k$, where each S_i is simple. The $*$ either switches between S_i and S_j , $i \neq j$, or preserves S_i . Both cases are referred to as $*$ -simple. The former case is called the exchange type. In the latter case, S_i is a simple algebra with an involution. Over any field, Wedderburn's theory gives a characterization of such simple algebras (see e.g. [AB95, Chap. 5]). Based on this, involutions on simple algebras are also classified [Alb39, Chap. X.4], and explicit lists for \mathbb{F}_q , \mathbb{R} , and \mathbb{C} can be found in [IQ19].

Given this structure, the idea is to reduce the search for a hyperbolic idempotent from general $*$ -algebras to simple $*$ -algebras. Clearly, if \mathfrak{A} contains a hyperbolic idempotent, then $\mathfrak{A}/\text{rad}(\mathfrak{A})$, and each $*$ -simple summand of $\mathfrak{A}/\text{rad}(\mathfrak{A})$, all contain hyperbolic idempotents. On the other hand, suppose each $*$ -simple summand of $\mathfrak{A}/\text{rad}(\mathfrak{A})$ contains a hyperbolic idempotent. Then the sum of these idempotents is a hyperbolic idempotent for $\mathfrak{A}/\text{rad}(\mathfrak{A})$. From here, to obtain a hyperbolic idempotent for \mathfrak{A} , we can use the classical idempotent lifting technique (see e.g. [Wil09b, Lemma 5.10]).

Therefore, it remains to handle the $*$ -simple case. Let \mathbb{K} denote some appropriate division algebra containing \mathbb{F} . The reader may as well think of \mathbb{K} as an extension field, as for \mathbb{F}_q , \mathbb{R} , and \mathbb{C} , the only “non-field” case is the quaternion algebra over \mathbb{R} . The exchange type is easy to handle: it is $*$ -isomorphic to $M(\ell, \mathbb{K}) \oplus M(\ell, \mathbb{K})^{op}$ with $(A, B)^* = (B, A)$. So one hyperbolic idempotent can be $(I, \mathbf{0})$. The simple case is more interesting. It is isomorphic to $M(\ell, \mathbb{K})$ with the involution defined by some non-degenerate classical form F ; this includes alternating, symmetric, and Hermitian ones.⁸ Then for $A \in M(\ell, \mathbb{K})$, $A^* = F^{-1}A^\dagger F$, where \dagger denotes either transpose (for alternating and symmetric) or conjugate transpose (for Hermitian). The problem is then to find a hyperbolic idempotent, or equivalently, an isotropic 2-decomposition, for this form F . But now this is a single form, so one can bring it to say a canonical form, and examine case by case.

For example, over \mathbb{C} , there are two types, symmetric and alternating. (The Hermitian type does not appear because $*$ is required to preserve \mathbb{C} .) A non-degenerate alternating form can always be transformed to $\begin{bmatrix} \mathbf{0} & I \\ -I & \mathbf{0} \end{bmatrix}$, so isotropic 2-decomposable. A non-degenerate symmetric form can be transformed to the identity matrix I . When I is of odd size ℓ , then it does not admit an isotropic 2-decomposition. Because if so, then I is isometric to $J = \begin{bmatrix} \mathbf{0} & A \\ A^t & \mathbf{0} \end{bmatrix}$, where A is of size $i \times (\ell - i)$. But then $\text{rk}(J)$ is either $2i$ or $2(\ell - i)$, an even number, so J is degenerate, a contradiction. When I is of even size ℓ , then it has an isotropic 2-decomposition. This is because we have $\begin{bmatrix} 1 & i \\ i & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & i \\ i & 1 \end{bmatrix} = \begin{bmatrix} 0 & 2i \\ 2i & 0 \end{bmatrix}$. We can then use this to bring I to $\begin{bmatrix} \mathbf{0} & 2iI' \\ 2iI' & \mathbf{0} \end{bmatrix}$, where I' is the $\ell/2 \times \ell/2$ identity matrix. Similar reasoning can be carried over \mathbb{R} and \mathbb{F}_q .

To make the above procedure constructive, we need to compute the algebra structure efficiently. This can be done, over \mathbb{F}_q with randomness [Rón90] and over \mathbb{C} deterministically [Ebe91, Rón94]. We also need to compute the $*$ -algebra structure, e.g. the forms associated with a simple $*$ -algebra, by [Wil09a, BW12, IQ19]. Finally, we need to compute the canonical forms of various forms by [Wil13, vdW05].

For the algorithm analysis, we work in the exact model, as going to extension fields is already unavoidable in computing the algebra structure. Over \mathbb{F}_q , one can always recover, from the solutions in the simple cases, an explicit hyperbolic projection matrix P over the original field, and then we

⁸While in principle this is correct, depending on the field, some type may not exist. For details see [IQ19].

can obtain the bases of the two subspaces in an isotropic 2-decomposition by computing the image and kernel of P . Over \mathbb{R} and \mathbb{C} , one can represent this projection matrix as a product of matrices over different extension fields [IQ19, Sec. 3.5].

This concludes an exposition of the algorithm.

7 Proof of Theorem 8

Let us first recall the alternative definition of non-commutative rank for a slightly more general situation. Given $\mathcal{B} \leq M(s \times t, \mathbb{F})$, an isotropic pair is a pair of vector spaces (U, V) , $U \leq \mathbb{F}^s$, $V \leq \mathbb{F}^t$, such that for any $u \in U, v \in V$, and any $B \in \mathcal{B}$, we have $u^t B v = 0$. The non-commutative rank of \mathcal{B} is then defined as $\text{ncrk}(\mathcal{B}) := (s + t) - \max\{c + e : c = \dim(U), e = \dim(V), (U, V) \text{ is an isotropic pair of } \mathcal{B}\}$. Note that the recent works [GGdOW16, IQS18] mostly deal with the setting that $s = t$. Suppose $\text{ncrk}(\mathcal{B}) = r$. By equivalence transformations, we can assume that every B is of the form $\begin{bmatrix} B_1 & B_2 \\ \mathbf{0} & B_3 \end{bmatrix}$, where B_2 is of size $a \times b$ such that $a + b = r$.

We review the setting for Theorem 8. Let $\mathcal{A} \leq \Lambda(n, \mathbb{F})$ be a bipartite alternating matrix space. By isometric transformations, we can assume that every $A \in \mathcal{A}$ is of the form $\begin{bmatrix} \mathbf{0} & B \\ -B^t & \mathbf{0} \end{bmatrix}$ where $B \in M(s \times t, \mathbb{F})$, $s + t = n$. All such B form a matrix space $\mathcal{B} \leq M(s \times t, \mathbb{F})$. We call such \mathcal{B} a matrix space induced from the bipartite structure of \mathcal{A} .

Before proving Theorem 8, let us examine some examples of isotropic spaces of \mathcal{A} .

1. First note that $\alpha(\mathcal{A}) \geq \max\{s, t\}$.
2. Second, suppose $\text{ncrk}(\mathcal{B}) = r$, so there exists $P \in GL(s, \mathbb{F})$ and $Q \in GL(t, \mathbb{F})$, such that every matrix in $P\mathcal{B}Q$ is of the form $\begin{bmatrix} B_1 & B_2 \\ \mathbf{0} & B_3 \end{bmatrix}$, where B_2 is of size $a \times b$ such that $a + b = r$. Let $R = \begin{bmatrix} P^t & \mathbf{0} \\ \mathbf{0} & Q \end{bmatrix}$. Then we have

$$\begin{aligned} R^t A R &= \begin{bmatrix} P & \mathbf{0} \\ \mathbf{0} & Q^t \end{bmatrix} \begin{bmatrix} \mathbf{0} & B \\ -B^t & \mathbf{0} \end{bmatrix} \begin{bmatrix} P^t & \mathbf{0} \\ \mathbf{0} & Q \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{0} & PBQ \\ -(PBQ)^t & \mathbf{0} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{0} & \mathbf{0} & B_1 & B_2 \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & B_3 \\ -B_1^t & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ -B_2^t & -B_3^t & \mathbf{0} & \mathbf{0} \end{bmatrix}, \end{aligned}$$

from which we get an isotropic space, consisting of the zero blocks in the middle part, of dimension $n - r$. (Note that the zero matrix at the (2, 3) position is of size $(s - a) \times (t - b)$, and the isotropic space corresponding to the zero blocks in the middle part is of size $(s - a) + (t - b) = (s + t) - (a + b) = n - r$.)

3. The third example we now describe represents a difference from the graph-theoretic setting. Suppose $s = t$, and every $B \in M(s, \mathbb{F})$ is symmetric. So $n = 2s$. Then let $U = \{u \in \mathbb{F}^n =$

$\mathbb{F}^{2s} : u = (u_1, \dots, u_s, u_1, \dots, u_s)^t \in \mathbb{F}^{2s} \leq \mathbb{F}^{2s}$. That is, U consists of those vectors whose i th component equals the $(i + s)$ th component, for $i \in [s]$, and $\dim(U) = s$. We claim that U is an isotropic space of \mathcal{A} . To see this, for a given $u \in U$, we can write it as $\begin{bmatrix} v \\ v \end{bmatrix}$ where $v \in \mathbb{F}^s$.

So for $A \in \mathcal{A}$ such that $A = \begin{bmatrix} \mathbf{0} & B \\ -B^t & \mathbf{0} \end{bmatrix}$, we have $\begin{bmatrix} v^t & v^t \end{bmatrix} \begin{bmatrix} \mathbf{0} & B \\ -B^t & \mathbf{0} \end{bmatrix} \begin{bmatrix} v \\ v \end{bmatrix} = -v^t B^t v + v^t B v = -v^t B v + v^t B v = 0$.

The proof of Theorem 8 basically suggests that isotropic spaces of the third type are not going to matter for the comparison with $\alpha(\mathcal{A})$. We now go into the proof.

Proof of Theorem 8. Let $r = \text{ncrk}(\mathcal{B})$ and $d = \alpha(\mathcal{A})$.

We first show that $\alpha(\mathcal{A}) \geq n - \text{ncrk}(\mathcal{B})$. Note that $r \leq \min\{s, t\}$, because we have the trivial isotropic pairs $(\mathbb{F}^s, \mathbf{0})$ and $(\mathbf{0}, \mathbb{F}^t)$. If $r = \min\{s, t\}$, note that $n - \min\{s, t\} = \max\{n - s, n - t\} = \max\{s, t\}$, and we do have isotropic spaces of dimensions s and t , respectively, by (1) from above. If $r < \min\{s, t\}$, then by (2) from above, there is an isotropic space of dimensions $n - r$. This shows that $\alpha(\mathcal{A}) \geq n - \text{ncrk}(\mathcal{B})$.

We then show that $\alpha(\mathcal{A}) \leq n - \text{ncrk}(\mathcal{B})$, or equivalently, $\text{ncrk}(\mathcal{B}) \leq n - \alpha(\mathcal{A})$. Again, if $\alpha(\mathcal{A}) = \max\{s, t\}$, then $\text{ncrk}(\mathcal{B}) \leq n - \max\{s, t\} = \min\{s, t\}$, which trivially holds. So we assume that $\alpha(\mathcal{A}) > \max\{s, t\}$. Let U be an isotropic space of dimension $d = \alpha(\mathcal{A})$, and take an $n \times d$ matrix whose columns form a basis of U , which, by abuse of notation, is also denoted by U . Let $U = \begin{bmatrix} v'_1 & v'_2 & \dots & v'_d \\ w'_1 & w'_2 & \dots & w'_d \end{bmatrix}$, where $v'_i \in \mathbb{F}^s$, and $w'_i \in \mathbb{F}^t$. Let $V' = [v'_1 \ v'_2 \ \dots \ v'_d]$, and $W' = [w'_1 \ w'_2 \ \dots \ w'_d]$. By doing a linear combination over the columns, we can assume that U is of the form $\begin{bmatrix} V & \mathbf{0} \\ W'' & W \end{bmatrix}$ where $V \in M(s \times c, \mathbb{F})$, $W'' \in M(t \times c, \mathbb{F})$, and $W \in M(t \times e, \mathbb{F})$, such that $c + e = d$, $\text{rk}(V) = c$, and $\text{rk}(W) = e$. By abuse of notation, let V be the subspace of \mathbb{F}^s spanned by columns in V , and W the subspace of \mathbb{F}^t spanned by columns in W .

Claim 23. *Let V and W be as above. Then (V, W) form an isotropic pair for \mathcal{B} .*

Proof. From above we have that $U = \begin{bmatrix} V & \mathbf{0} \\ W'' & W \end{bmatrix}$, and suppose $U = \begin{bmatrix} v_1 & v_2 & \dots & v_d \\ w_1 & w_2 & \dots & w_d \end{bmatrix}$. This means that $v_i = \mathbf{0}$ for $c < i \leq d$. Since U is an isotropic space of \mathcal{A} , we have, for any $i, j \in [d]$, and $A = \begin{bmatrix} \mathbf{0} & B \\ -B^t & \mathbf{0} \end{bmatrix} \in \mathcal{A}$, $\begin{bmatrix} v_i^t & w_i^t \end{bmatrix} \begin{bmatrix} \mathbf{0} & B \\ -B^t & \mathbf{0} \end{bmatrix} \begin{bmatrix} v_j \\ w_j \end{bmatrix} = -w_i^t B^t v_j + v_i^t B w_j = -v_j^t B w_i + v_i^t B w_j = 0$, so $v_i^t B w_j = v_j^t B w_i$. In particular, for any column vector $v \in V$ and any column vector $w \in W$, we have $v^t B w = \mathbf{0}^t B w'' = 0$ for some column vector $w'' \in W''$. This concludes the proof. \square

We then get that $\text{ncrk}(\mathcal{B}) \leq (s + t) - (c + e) = n - d = n - \alpha(\mathcal{A})$, and the proof is concluded. \square

We now set out to prove Corollary 9. For this we need one more ingredient. In the literature [GGdOW16, IQS18], the computation of the non-commutative ranks only deals with the case when $s = t$. To use that for general $M(s \times t, \mathbb{F})$ we need a little twist.

Proposition 24. *Over any field \mathbb{F} , computing the non-commutative rank of $\mathcal{B} \leq M(s \times t, \mathbb{F})$ can be done in deterministic polynomial time.*

Proof. If $s = t$, this follows from [IQS18]. Without loss of generality, let us assume then $s < t$. We shall construct some $\mathcal{C} \leq \mathbb{M}(t, \mathbb{F})$ as follows. First, for any $B \in \mathcal{B}$, let $B' = \begin{bmatrix} \mathbf{0} \\ B \end{bmatrix}$, where $\mathbf{0}$ is of size $(t-s) \times t$, so $B' \in \mathbb{M}(t, \mathbb{F})$. Second, recall that $E_{i,j}$ is the elementary matrix with the (i, j) th entry being 1, and the rest entries being 0. Then \mathcal{C} is the matrix space spanned by all B' and $E_{i,j}$ with $1 \leq i \leq t-s$ and $1 \leq j \leq t$.

We claim that $\text{ncrk}(\mathcal{B}) + (t-s) = \text{ncrk}(\mathcal{C})$. To see that $\text{ncrk}(\mathcal{B}) + (t-s) \geq \text{ncrk}(\mathcal{C})$, let (U, V) be an isotropic pair of \mathcal{B} , where $U \leq \mathbb{F}^s$, $V \leq \mathbb{F}^t$, such that $\text{ncrk}(\mathcal{B}) = s+t - \dim(U) - \dim(V)$. Let $U' \leq \mathbb{F}^t$ be the image of U under the embedding \mathbb{F}^s to \mathbb{F}^t by sending e_i to e_{i+t-s} . Clearly, (U', V) is an isotropic pair for \mathcal{C} , so $\text{ncrk}(\mathcal{C}) \leq 2t - \dim(U') - \dim(V) = (t-s) + s+t - \dim(U) - \dim(V) = (t-s) + \text{ncrk}(\mathcal{B})$.

To show that $\text{ncrk}(\mathcal{B}) + (t-s) \leq \text{ncrk}(\mathcal{C})$, let (U, V) be an isotropic pair of \mathcal{C} . If $\text{ncrk}(\mathcal{C}) = t$, then the equality is trivial. So in the following we assume $\text{ncrk}(\mathcal{C}) < t$. We claim that if $V \neq \mathbf{0}$, then U is a subspace of $\langle e_{t-s+1}, \dots, e_t \rangle$. Suppose not, then U contains a vector $u = [u_1, \dots, u_t]^t$ with some $u_i \neq 0$ for $1 \leq i \leq t-s$. Because $E_{i,j}$ is present in \mathcal{C} for $1 \leq j \leq t$, for $v \in \mathbb{F}^n$ to satisfy that $u^t E_{i,j} v = 0$ for any $1 \leq j \leq t$, v has to be $\mathbf{0}$. This implies that V has to be $\mathbf{0}$. Therefore all isotropic pairs of \mathcal{C} , except a trivial one $(\mathbb{F}^t, \mathbf{0})$, are also isotropic pairs of \mathcal{B} . Therefore, if $\text{ncrk}(\mathcal{C}) < t$, and (U, V) is an isotropic pair for $\text{ncrk}(\mathcal{C})$ such that $\text{ncrk}(\mathcal{C}) = 2t - \dim(U) - \dim(V) < t$, we have $V \neq \mathbf{0}$, so $U \leq \langle e_{t-s+1}, \dots, e_t \rangle$. Let U' be the image of U under the projection from \mathbb{F}^t to \mathbb{F}^s by sending $(v_1, \dots, v_t)^t$ to $(v_{t-s+1}, \dots, v_t)^t$. We see then $\dim(U') = \dim(U)$, and (U', V) is an isotropic space for \mathcal{B} . From this we can conclude the proof. \square

We are now ready to prove Corollary 9.

Proof of Corollary 9. Given $\mathcal{A} \leq \Lambda(n, \mathbb{F}_q)$, q odd, first put it into the explicit bipartite form using Theorem 6, which also produces the bases of the two subspaces in an isotropic 2-decomposition. Then for a matrix space $\mathcal{B} \leq \mathbb{M}(s \times t, \mathbb{F})$ induced from the bipartite structure, compute its non-commutative rank $r = \text{ncrk}(\mathcal{B})$ using Proposition 24. The isotropic number of \mathcal{A} is then $n - r$, by Theorem 8. \square

Remark 25. To obtain analogues of Corollary 9 over \mathbb{R} and \mathbb{C} , the bottleneck is that over \mathbb{R} and \mathbb{C} , Theorem 6 only outputs the projection matrix as the product of a sequence of matrices over different extension fields. This prevents us from working with the bases of the subspaces in an isotropic 2-decomposition directly. Of course, if we are content with approximating those algebraic numbers up to certain precision, we can use the representation of the projection as a product of matrices over different extension fields to get such, and then work with that.

8 Proof of Theorem 10

Let us first work with general \mathbb{F} , and then restrict to our target field \mathbb{Q} at some point. Recall that the problem is to decide whether an alternating matrix space $\mathcal{A} \leq \Lambda(n, \mathbb{F})$ has an isotropic space of dimension 2. We first make the following easy observation.

Observation 26. *Let $\mathcal{A} = \langle A_1, \dots, A_m \rangle \leq \Lambda(n, \mathbb{F})$. Then \mathcal{A} has an isotropic space of dimension 2, if and only if there exist linearly independent $v, w \in \mathbb{F}^n$ such that for any $A \in \mathcal{A}$, $v^t A w = 0$, which is further equivalent to that for any $i \in [m]$, $v^t A_i w = 0$.*

We now need to prove some auxiliary results. Let $\mathcal{B} = \langle B_1, \dots, B_m \rangle \leq M(n, \mathbb{F})$, and let $\mathbf{B} = (B_1, \dots, B_m) \in M(n, \mathbb{F})^m$. Here is a natural problem about matrix spaces.

Problem 1. The existential singularity problem for matrix spaces, or the linear \exists -singularity problem, asks the following: given $\mathcal{B} \in M(n, \mathbb{F})$, decide whether there exists a singular (e.g. non-full-rank) *non-zero* matrix in \mathcal{B} .

The linear \exists -singularity problem turns out to be quite interesting. We discuss this problem in detail in Section 8.1. For the sake of proving Theorem 10, we need the following result, whose proof can be found there.

Lemma 27. *Over \mathbb{Q} , assuming the generalized Riemann hypothesis, there is a randomized polynomial-time reduction from deciding quadratic residuosity modulo squarefree composite numbers to the linear \exists -singularity problem.*

We then show that the linear \exists -singularity problem reduces to deciding whether an alternating matrix space has an isotropic space of dimension 2. This reduction works over any field.

For this purpose, it will be convenient to define an intermediate problem, which may be viewed as just a reformulation of Problem 1.

Recall that $\mathcal{B} = \langle B_1, \dots, B_m \rangle \leq M(n, \mathbb{F})$, and $\mathbf{B} = (B_1, \dots, B_m) \in M(n, \mathbb{F})^m$. Think of \mathbf{B} as a 3-tensor $T_{\mathbf{B}}$ of size $n \times n \times m$, such that $T_{\mathbf{B}}(i, j, k) = B_k(i, j)$. That is, B_k 's are the slices according to the third index (lateral slices). We will also be interested in the matrices obtained according to the first index (horizontal slices) and the second index (vertical slices). Specifically, define \mathbf{B}^v be the n -tuple of $n \times m$ matrices that are vertical slices of $T_{\mathbf{B}}$. That is, $\mathbf{B}' = (B'_1, \dots, B'_n) \in M(n \times m, \mathbb{F})^n$, where $B'_j = [B_1 e_j, \dots, B_m e_j]$, or in other words, $B'_j(i, k) = T_{\mathbf{B}}(i, j, k) = B_k(i, j)$. Similarly we can define the matrix tuple consisting of the horizontal slices of $T_{\mathbf{B}}$.

We now consider the matrix space $\mathbf{B}' \in M(n \times m, \mathbb{F})^n$. For $v = (v_1, \dots, v_m)^t \in \mathbb{F}^m$, its right degree in \mathbf{B}' is defined to be the rank of $[B'_1 v, \dots, B'_n v] = v_1 B_1 + \dots + v_m B_m$. Therefore, every non-zero $v \in \mathbb{F}^m$ has right degree n in \mathbf{B}' , if and only if every matrix in \mathcal{B} is of rank n . Lemma 27 then immediately implies the following.

Corollary 28. *Over \mathbb{Q} , assuming the generalized Riemann hypothesis, there is a randomized polynomial-time reduction from deciding quadratic residuosity modulo squarefree composite numbers to deciding whether there exists a non-zero $v \in \mathbb{F}^m$ of right degree $< n$ w.r.t. a matrix tuple $\mathbf{B}' \in M(n \times m, \mathbb{F})$.*

Given $\mathbf{B}' = (B'_1, \dots, B'_n) \in M(n \times m, \mathbb{F})$, we construct a tuple of alternating matrices of size $(n + m) \times (n + m)$, as follows. For $i \in [n]$, let $A_i = \begin{bmatrix} \mathbf{0} & B'_i \\ -B_i{}^t & \mathbf{0} \end{bmatrix}$. For $1 \leq i < j \leq n$, let $C_{i,j} = e_i e_j^t - e_j e_i^t$. For $1 \leq k < \ell \leq m$, let $D_{k,\ell} = e_{n+k} e_{n+\ell}^t - e_{n+\ell} e_{n+k}^t$. Note that $C_{i,j}$ and $D_{k,\ell}$ are elementary alternating matrices. Let $\mathbf{A} = (A_1, \dots, A_n, C_{1,2}, \dots, C_{n-1,n}, D_{1,2}, \dots, D_{m-1,m})$. We now claim the following.

Claim 29. *Let $\mathbf{B}' \in M(n \times m, \mathbb{F})^n$ and $\mathbf{A} \in \Lambda(n + m, \mathbb{F})^{n + \binom{n}{2} + \binom{m}{2}}$ be as above. Let $\mathcal{A} = \langle \mathbf{A} \rangle \leq \Lambda(n + m, \mathbb{F})$. Then there exists a non-zero $v' \in \mathbb{F}^m$ of right degree $< n$ in \mathbf{B}' if and only if \mathcal{A} has an isotropic space of dimension 2.*

Proof. By Observation 26, to decide whether \mathcal{A} has an isotropic space of dimension ≥ 2 , we only need to test whether there exist linearly independent $u, v \in \mathbb{F}^{n+m}$, such that for any $E = A_i$, or $C_{i,j}$, or $D_{k,\ell}$, $u^t E v = 0$.

The only if direction is easy to verify. Suppose $v' \in \mathbb{F}^m$ is of right degree $n-1$ w.r.t. \mathbf{B}' , namely $[B'_1 v', \dots, B'_m v']$ is of rank $< n$. Then take any non-zero $u' \in \mathbb{F}^n$ in the left kernel of $[B'_1 v', \dots, B'_m v']$, and we have that for $i \in [m]$, $u'^t B'_i v' = 0$. Now construct $u = \begin{bmatrix} u' \\ \mathbf{0} \end{bmatrix} \in \mathbb{F}^{n+m}$, and $v = \begin{bmatrix} \mathbf{0} \\ v' \end{bmatrix} \in \mathbb{F}^{n+m}$. For $i \in [n+m]$, let the i th component of u (resp. v) be u_i (resp. v_i). Then for $i \in \{n+1, \dots, n+m\}$, $u_i = 0$. For $i \in [n]$, $v_i = 0$. Clearly, u and v are linearly independent. Furthermore, it is easy to verify that (1) $u^t A_i v = u'^t B'_i v' = 0$, (2) $u^t C_{i,j} v = u_i v_j - u_j v_i = u_i \cdot 0 - u_j \cdot 0 = 0$, as $i, j \in [n]$, and similarly (3) $u^t D_{k,\ell} v = 0$. Then u and v spans a dimension-2 isotropic space of \mathcal{A} .

For the if direction, suppose u and v are linearly independent vectors in \mathbb{F}^{n+m} , and satisfy that for any $E = A_i$, or $C_{i,j}$, or $D_{k,\ell}$, $u^t E v = 0$. Write $u = \begin{bmatrix} u_1 \\ u_2 \end{bmatrix}$, where $u_1 \in \mathbb{F}^n$ and $u_2 \in \mathbb{F}^m$. Similarly write $v = \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}$, where $v_1 \in \mathbb{F}^n$ and $v_2 \in \mathbb{F}^m$. By $u^t C_{i,j} v = 0$, we have that u_1 and v_1 are linearly dependent. By $u^t D_{k,\ell} v = 0$, we have that u_2 and v_2 are linearly dependent. We first observe that it cannot be the case that $u_1 = v_1 = \mathbf{0}$, nor $u_2 = v_2 = \mathbf{0}$. As otherwise, say if $u_1 = v_1 = \mathbf{0}$, then because u_2 and v_2 are linearly dependent, we have u and v are linearly dependent, a contradiction. Therefore, without loss of generality, we assume that $u_1 \neq \mathbf{0}$, so $v_1 = \alpha_1 u_1$ for some $\alpha_1 \in \mathbb{F}$. We then have two cases. In the first case, if $u_2 = \mathbf{0}$, then $v_2 \neq \mathbf{0}$, and $v' = v - \alpha_1 u$ and u are linearly independent. In the second case, if $u_2 \neq \mathbf{0}$, then $v_2 \neq \alpha_1 u_2$, as otherwise u and v would be linearly dependent. Then again, letting $v' = v - \alpha_1 u$, we have u and v' are linearly independent. Clearly, u and v' still satisfy that for any $E = A_i$, or $C_{i,j}$, or $D_{k,\ell}$, $u^t E v' = 0$. Write v' as $\begin{bmatrix} v'_1 \\ v'_2 \end{bmatrix}$ where $v'_1 \in \mathbb{F}^n$, and $v'_2 \in \mathbb{F}^m$, so $v'_1 \neq \mathbf{0}$, and $v'_2 \neq \mathbf{0}$. We then have $u_2 = \alpha_2 v'_2$. Letting $u' = u - \alpha_2 v'$, we have u' and v' are linearly independent, and for any $E = A_i$, or $C_{i,j}$, or $D_{k,\ell}$, $u'^t E v' = 0$. Write u' as $\begin{bmatrix} u'_1 \\ u'_2 \end{bmatrix}$ where $u'_1 \in \mathbb{F}^n$, and $u'_2 \in \mathbb{F}^m$, so $u'_1 \neq \mathbf{0}$, and $u'_2 = \mathbf{0}$. It is then straightforward to verify that the condition $u'^t A_i v' = 0$ is equivalent to $u'^t_1 B'_i v'_2 \neq 0$. Recall that neither u'_1 nor v'_2 is the zero vector; this just translates to say that v'_2 is of right degree $< n$ w.r.t. \mathbf{B}' . \square

Theorem 10 follows by combining Claim 29 and Corollary 28.

8.1 The existential singularity problem for matrix spaces

In this subsection, we discuss on Problem 1, which we believe is a very interesting problem in its own right. We therefore examine this problem over various fields, and prove Lemma 27 over \mathbb{Q} .

The affine version of Problem 1 has been studied in [BFS99]. More specifically, the \exists -singularity problem for affine matrix spaces asks to decide whether an affine matrix space contains a non-full-rank matrix (not necessarily non-zero). In [BFS99], this problem was called the singularity problem. This may cause some confusion, because in [GGdOW16, IQS18] the singularity problem for matrix spaces is to decide whether all matrices in a matrix space are singular. For clarification, we then call the problem in [BFS99] the affine \exists -singularity problem, and Problem 1 the linear \exists -singularity problem.

In [BFS99], it was shown that the affine \exists -singularity problem is NP-hard over \mathbb{F}_q , \mathbb{Q} , or \mathbb{R} . We first note that the linear \exists -singularity problem reduces to that for affine matrix spaces, but the inverse direction is not clear.⁹ Furthermore, the proof strategy of [BFS99] cannot be adapted

⁹To reduce the matrix space case to the affine case is easy: if $\mathcal{B} = \langle B_1, \dots, B_m \rangle$, then form m affine spaces

directly to tackle Problem 1, because of the introduction of field constants in the reduction. In particular, the use of field constants in the transformation from algebraic branching programs to symbolic determinants by Valiant [Val79] seems particularly crucial. Indeed, as we will see below, the \exists -singular problems for matrix spaces and for affine matrix spaces demonstrate quite different behaviors.

Matrix spaces in which every non-zero matrix is of full-rank has been studied in mathematics for a long time. More broadly, if $\mathcal{B} \leq M(n, \mathbb{F})$ satisfies that every non-zero matrix in \mathcal{B} is of a fixed rank r , we say that \mathcal{B} satisfies the fixed rank r condition. Such matrix spaces are of interests in algebraic geometry (mostly when over algebraically closed fields), differential topology (mostly when over \mathbb{R}), number theory (mostly when over \mathbb{Q}), and algebra (mostly when over finite fields). It turns out that several results from these different branches of mathematics will be useful for our algorithmic purposes as well.

To start with, the following quantity has been studied extensively in the literature. Let $\rho(n, r, \mathbb{F})$ be the maximum dimension over those $\mathcal{B} \leq M(n, \mathbb{F})$ satisfying the fixed rank r condition. Also let $\tau(n, r, \mathbb{F})$ be the maximum dimension over those affine matrix spaces $\mathcal{C} \subseteq M(n, \mathbb{F})$ satisfying the fixed rank r condition. Also let $\rho(n, \mathbb{F}) := \rho(n, n, \mathbb{F})$, and $\tau(n, \mathbb{F}) := \tau(n, n, \mathbb{F})$. As pointed out in [dSP13], $\rho(n, \mathbb{F}) \leq n$, and $\tau(n, \mathbb{F}) = \binom{n}{2}$. Two remarks are due here. First, $\rho(n, \mathbb{F})$ can be much smaller than n for certain fields; see below. Second, the $\binom{n}{2}$ bound for $\tau(n, \mathbb{F})$ can be easily achieved at $I_n + \mathcal{U}$ where \mathcal{U} is the linear space of strictly upper triangular matrices. This distinction already suggests that the difference between the linear and the affine cases can be significant.

In this following, we discuss on \mathbb{C} , \mathbb{R} , and \mathbb{Q} , comparing the linear and affine settings, and presenting some algorithms for the linear case, including a proof of Lemma 27. We refer the interested reader to [She11] for the finite field case.

Over \mathbb{C} . The affine \exists -singularity problem over \mathbb{C} is only known to be in RP [BFS99].

We then turn to the linear \exists -singularity problem over \mathbb{C} . Sylvester showed that $\rho(n, \mathbb{C}) \leq 1$ [Syl86], and Westwick generalized that to $\rho(n, r, \mathbb{C}) \leq 2n - 2r + 1$ [Wes87]. Some subsequent developments include [IL99, BFM13].

Sylvester’s result immediately translates to a deterministic efficient algorithm for the linear \exists -singularity problem over \mathbb{C} : if the input matrix space $\mathcal{B} \leq M(n, \mathbb{C})$ is of dimension ≥ 2 , then return “exists.” Otherwise, $\mathcal{B} = \langle B \rangle$, and return “exists” if and only if B is of full-rank

Over \mathbb{R} . The affine \exists -singularity problem over \mathbb{R} is NP-hard [BFS99].

We then turn to the linear \exists -singularity problem over \mathbb{R} . Based on the Radon-Hurwitz construction and Adams’ vector field theorem [Ada62], $\rho(n, \mathbb{R})$ is equal to the so-called Hurwitz-Radon function (see [ALP65a]). For $n \in \mathbb{N}$, write n in the form of $2^{4a+b} \cdot (2c+1)$ where $b \in \{0, 1, 2, 3\}$, and the Hurwitz-Radon function is $HR(n) = 8a + 2^b$. The significance of this result for our algorithmic purpose is that $HR(n) \leq 2(\log n + 4)$. Some subsequent developments include [Mes90, LY93, CP07].

Therefore, the linear \exists -singularity problem over \mathbb{R} admits the following quasipolynomial-time algorithm. If the input matrix space $\mathcal{B} \leq M(n, \mathbb{C})$ is of dimension $\geq HR(n)$, then return “exists.” Otherwise, $\mathcal{B} = \langle B_1, \dots, B_m \rangle$ where $m \leq HR(n) \leq 2(\log n + 4)$. We form m affine spaces, $\mathcal{C}_i := B_i + \langle B_1, \dots, B_{i-1}, B_{i+1}, \dots, B_m \rangle$, for every $i \in [m]$. The question then becomes whether any of the \mathcal{C}_i ’s contains a singular matrix. This can be done by computing $f_i := \det(B_1x_1 + \dots + B_{i-1}x_{i-1} + B_i + B_{i+1}x_{i+1} + \dots + B_mx_m)$ explicitly. Since the polynomial f_i , involving $O(\log n)$

$B_i + \langle B_1, \dots, B_{i-1}, B_{i+1}, \dots, B_m \rangle$.

variables, is of degree n , f_i has $n^{O(\log n)}$ monomials, and we can fully write out f_i in time polynomial in $n^{O(\log n)}$.¹⁰ After that, we can use the existential theory of reals [Can88, Ren92] to determine whether f_i has a non-trivial zero in time $n^{O(\log n)}$. Return “exists” if and only if one of these f_i ’s is solvable. This concludes the proof.

Over \mathbb{Q} . The affine \exists -singularity problem over \mathbb{Q} is NP-hard [BFS99].

We then turn to the linear \exists -singularity problem over \mathbb{Q} . To start with, observe that $\rho(n, \mathbb{Q}) \geq n$. This is because we can take a degree- n extension field \mathbb{K} over \mathbb{Q} , and use the regular representation of \mathbb{K} . We now prove Lemma 27, which suggests that the linear \exists -singularity problem is not so easy either.

Proof of Lemma 27. We consider a special case of Problem 1 as follows. Assume $\mathcal{B} \leq M(n, \mathbb{Q})$ is closed under matrix multiplication, so \mathcal{B} forms an algebra over \mathbb{Q} . In this setting, Problem 1 just asks whether \mathcal{B} is not a division algebra. We can even specialize further by considering \mathcal{B} being a central simple algebra over \mathbb{Q} .

In [Rón87], Rónyai considered the problem of testing whether a central simple algebra over \mathbb{Q} of dimension 4 is isomorphic to $M(2, \mathbb{Q})$. He showed that assuming the generalized Riemann hypothesis, there is a randomized efficient reduction from deciding quadratic residuosity modulo squarefree composite numbers to this problem. In [Rón87], the algebras are represented by structural constants, but these can be turned into matrix representations in $M(4, \mathbb{Q})$ (see e.g. [IR99]). It follows that there is an analogous reduction for matrix algebras.

We can then conclude the proof, because such an algebra is either isomorphic to $M(2, \mathbb{Q})$ (in which there exists a non-zero singular matrix) or a division algebra (in which every non-zero matrix is full-rank). \square

9 Proof of Theorem 11

9.1 Some basic statistics

All results in this subsection are either classical or straightforward. We collect them here, and provide proofs, partly for completeness, and partly because we will use some of the arguments here in the following.

We first recall the following bound on the number of subspaces of \mathbb{F}_q^n .

Fact 30. 1. For $d \leq \mathbb{N}$, $0 \leq d \leq n$, the number of dimension- d subspaces of \mathbb{F}_q^n is equal to the Gaussian binomial coefficient

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

2. The Gaussian binomial coefficient satisfies:

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

¹⁰There are several ways of doing this. One approach is to transform the determinant expression into an arithmetic circuit, and then compute along this circuit to get the final polynomial.

3. The number of subspaces of \mathbb{F}_q^n is $q^{\frac{1}{4}n^2 + \Theta(n)}$.

Proof. (1) is well-known. For (2), it is enough to verify that for any prime power q , and $n, d, k \in \mathbb{N}$, $n \geq d > k$, we have

$$q^{n-d} \leq \frac{q^n - q^k}{q^d - q^k} \leq q^{n-d+1}.$$

For (3), it is well-known that $\binom{n}{d}_q$ achieves maximal over d at $d = \lfloor n/2 \rfloor$. So we have

$$q^{\frac{1}{4}n^2 - \frac{1}{4}} \leq \binom{n}{\lfloor n/2 \rfloor}_q \leq \sum_{d=0}^n \binom{n}{d}_q \leq (n+1) \cdot \binom{n}{\lfloor n/2 \rfloor}_q \leq q^{\frac{1}{4}n^2 + \lfloor n/2 \rfloor + \log(n+1)},$$

from which the result follows. \square

Analogously, we consider the number of isotropic spaces of a non-degenerate alternating form $A \in \Lambda(n, q)$.

Proposition 31. *Let $A \in \Lambda(n, q)$, n even, be a non-degenerate alternating matrix. Then we have the following.*

1. For $d \in \mathbb{N}$, $0 \leq d \leq n/2$, the number of dimension- d isotropic spaces of A is

$$I(A, d) := \frac{(q^n - 1) \cdot (q^{n-1} - q) \cdot \dots \cdot (q^{n-(d-1)} - q^{d-1})}{(q^d - 1) \cdot (q^d - q) \cdot \dots \cdot (q^d - q^{d-1})}.$$

For $d \in \mathbb{N}$, $d > n/2$, there are no dimension- d isotropic spaces.

2. For $d \in \mathbb{N}$, $0 \leq d \leq n/2$, $I(A, d)$ is bounded as follows:

$$q^{nd - \frac{3}{2}d^2 + \frac{1}{2}d} \leq I(A, d) \leq q^{nd - \frac{3}{2}d^2 + \frac{3}{2}d}.$$

3. The number of isotropic spaces of A is $q^{\frac{1}{6}n^2 + \Theta(n)}$.

4. The number of maximal isotropic spaces of A is $q^{\frac{1}{8}n^2 + \Theta(n)}$.

Proof. For (1), suppose we have chosen u_1, \dots, u_i such that $\langle u_1, \dots, u_i \rangle$ is an isotropic space. We then need to select the next eligible u_{i+1} , such that $\langle u_1, \dots, u_{i+1} \rangle$ forms an isotropic space. Since u_{i+1} needs to satisfy $u_{i+1}^t A u_j = 0$ for $1 \leq j \leq i$, and A is non-degenerate, u_{i+1} should be from a dimension- $(n-i)$ subspace, namely the subspace orthogonal to Au_j , $1 \leq j \leq i$. Furthermore, u_{i+1} is not in $\langle u_1, \dots, u_i \rangle$. So there are $q^{n-i} - q^i$ choices of u_{i+1} in the i th step. This explains the numerator. The denominator is of such form, because for each isotropic space there are these many ordered bases.

For (2), it follows from the same argument as the proof for Fact 30 (2).

For (3), note that $nd - \frac{3}{2}d^2$ achieves its maximum at $d = \frac{1}{3}n$.

For (4), note that maximal isotropic spaces are of dimension $n/2$. This is because for any isotropic space U of dimension $d < n/2$, we can choose an eligible u_{d+1} as in the proof for (1), such that $\langle U, u_{d+1} \rangle$ is also an isotropic space. \square

Proposition 32. *Let $A \in \Lambda(n, q)$, n even, be a non-degenerate alternating matrix. Then all isotropic spaces of A can be enumerated in time $q^{\frac{1}{6}n^2 + O(n)}$.*

Proof. We enumerate isotropic spaces according to dimensions in an increasing order. Each subspace of \mathbb{F}_q^n is represented by an ordered basis. We will maintain a list L of all isotropic subspaces, and for each isotropic space U of dimension d , maintain a list of isotropic spaces of dimension $d+1$ that contain U , denoted as $L(U)$. Note that for a fixed U , there are at most q^{n-d} such spaces. In other words, we will record the lattice of isotropic spaces.

Suppose we have enumerated all isotropic spaces of dimensions $\leq d$. To enumerate isotropic spaces of dimension $d+1$, we maintain a list of such spaces. Then for each isotropic space U of dimension d , and for each $u \in \text{rad}(U) \setminus U$, we form $U' = \langle u, U \rangle$, and test whether U' is in $L(U)$. If not, then we add U' to L . We also add it to $L(U)$, and for every dimension d -subspace \tilde{U} of U' , add U' to $L(\tilde{U})$. Otherwise we move on.

Clearly, in the above procedure, each isotropic space will be added, and only added to L once. This procedure runs in time $N \cdot q^{O(n)}$, where N denotes the number of isotropic spaces of A . We can then conclude by resorting to Proposition 31 (3). \square

When working with maximal isotropic spaces, it is enough to restrict our attention to just non-degenerate matrix spaces.

Observation 33. For $\mathcal{A} \leq \Lambda(n, \mathbb{F})$, any maximal isotropic space of \mathcal{A} contains $\text{rad}(\mathcal{A})$.

9.2 A non-trivial upper bound on the number of maximal isotropic spaces

For $\mathcal{A} \leq \Lambda(n, q)$, let $\text{MI}(\mathcal{A})$ be the set of maximal isotropic spaces of \mathcal{A} , and $\text{NMI}(\mathcal{A})$ be the number of maximal isotropic spaces of \mathcal{A} , e.g. the size of $\text{MI}(\mathcal{A})$. Let $\text{MaxNMI}(n, q)$ be the maximum of $\text{NMI}(\mathcal{A})$ over all $\mathcal{A} \leq \Lambda(n, q)$. By Fact 30 (3) and Proposition 31 (4),

$$q^{\frac{1}{4}n^2 + O(n)} \geq \text{MaxNMI}(n, q) \geq q^{\frac{1}{8}n^2 + \Omega(n)}.$$

Theorem 11, slightly reformulated. Let $\text{MaxNMI}(n, q)$ be as above. Then $\text{MaxNMI}(n, q) \leq q^{\frac{1}{6}n^2 + C \cdot n}$ for some large enough absolute constant C .

Let us illustrate the proof strategy for Theorem 11, before we enter the details.

The starting point of our proof is the alternative proof bounding the number of maximal independent sets by Wood [Woo11].

The core of Wood's argument is the following. Let $G = (V, E)$ be a graph on n vertices. Recall that we want to prove that the number of maximal independent sets in G is no more than $g(n) = 3^{\frac{n}{3}}$. We shall do an induction on n . Let $v \in V$ be a vertex of minimal degree d . Let $N(v)$ be the set of neighbours of v together with v (e.g. the closed neighbourhood of v). Then any maximal independent set I contains some $w \in N(v)$, as otherwise $I \cup \{v\}$ would be a larger independent set. If I contains $w \in N(v)$, then $I \setminus \{w\}$ would be a maximal independent set of $G|_{V \setminus N(w)}$, the induced subgraph of G on $V \setminus N(w)$. Since $G|_{V \setminus N(w)}$ is of size $\leq n - d - 1$, we then have

$$g(n) \leq (d+1) \cdot g(n-d-1). \tag{3}$$

From this relation and the induction hypothesis, the result follows in a rather straightforward fashion.

In the following, we will develop a linear algebraic analogue of Equation 3. However, just applying this does not suffice, when there are many vectors of degree 1.

We remedy this by showing that in this setting, the maximum rank is large, which allows us to use an argument similar to one in Proposition 31. More specifically, recall that in Proposition 31 (3), we showed that the number of isotropic spaces of a non-degenerate alternating matrix is bounded from above by $q^{\frac{1}{6}n^2+O(n)}$. Note that any maximal isotropic space of \mathcal{A} is an isotropic space of any $A \in \mathcal{A}$. So if \mathcal{A} contains a non-degenerate A , we can immediately obtain Theorem 11 in this case. However, there are non-degenerate matrix spaces that do not contain non-degenerate alternating matrices. For example, the following is an alternating matrix space of maximum rank 2, written in a parametrized form:

$$A = \begin{bmatrix} 0 & x_1 & \dots & x_n \\ -x_1 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ -x_n & 0 & \dots & 0 \end{bmatrix}.$$

For our purpose, we will need to bound the number of isotropic spaces for matrix spaces of rank $> \frac{2}{3}n$. So the following lemma is required; its proof is postponed to Section 9.2.1.

Lemma 34. *Let $A \in \Lambda(n, q)$ be of rank $> \frac{2}{3}n$. Then the number of isotropic spaces of A is bounded from above by $q^{\frac{1}{6}n^2+Dn}$ for some large enough absolute constant D .*

We are now ready to prove Theorem 11.

Proof of Theorem 11. Let $\mathcal{A} \leq \Lambda(n, q)$. By Observation 33, for our purpose, we can assume that \mathcal{A} is non-degenerate. Let $g_q(n) = q^{\frac{1}{6}n^2+Cn}$. We prove by an induction on n . Assume $\text{MaxNMI}(\ell, q) \leq g_q(\ell)$ holds for any $\ell < n$. Our goal is to show that $\text{NMI}(\mathcal{A}) \leq g_q(n)$.

Let $d = \min\{\deg_{\mathcal{A}}(v) : v \in \mathbb{F}_q^n, v \neq \mathbf{0}\}$. As \mathcal{A} is non-degenerate, $d \geq 1$. Take any $v \in \mathbb{F}_q^n$ of degree d , and let $c = n - d$ be the codegree of v . Let $N(v) := (\mathbb{F}_q^n \setminus \text{rad}_{\mathcal{A}}(v)) \cup \{v\} = \{u \in \mathbb{F}_q^n : \exists A \in \mathcal{A}, u^t A v \neq 0\} \cup \{v\}$. We call $N(v)$ the closed neighbourhood of v . Note that $|N(v)| = q^n - q^c + 1$.

Let $U \leq \mathbb{F}_q^n$ be a maximal isotropic space of \mathcal{A} . Clearly, $U \cap N(v) \neq \emptyset$. As otherwise, we have $U \subseteq \text{rad}(v)$ and $v \notin U$. This is equivalent to that $v \in \text{rad}(U)$ and $v \notin U$. It follows that $U \subsetneq \text{rad}(U)$, so by Observation 19, U is not maximal, a contradiction.

Therefore there exists some $w \in N(v) \cap U$. It follows that $U \subseteq \text{rad}(w)$. Since U is maximal isotropic in \mathcal{A} , U is also a maximal isotropic space of $\mathcal{A}|_{\text{rad}(w)}$. As $\deg(w) \geq \deg(v) = d$, $\dim(\text{rad}(w)) \leq c$. Furthermore, note w is an isolated vector in $\mathcal{A}|_{\text{rad}(w)}$. We then have

$$\text{NMI}(\mathcal{A}) \leq \sum_{w \in N(v)} \text{NMI}(\mathcal{A}|_{\text{rad}(w)}) \tag{4}$$

$$\leq (q^n - q^c + 1) \cdot g_q(c - 1), \tag{5}$$

where the second inequality is due to the induction hypothesis. Note that on the right hand side, we have $g_q(c - 1)$ instead of $g_q(c)$, because w is an isolated vector in $\mathcal{A}|_{\text{rad}(w)}$, and Observation 33. The reader may want to compare this with Equation 3.

Now suppose $d \geq 2$, that is, $c \leq n - 2$. We then have

$$\begin{aligned} \text{NMI}(\mathcal{A}) &\leq q^n \cdot g_q(n - 3) \\ &\leq q^n \cdot q^{\frac{1}{6}(n-3)^2+C(n-3)} \\ &= q^{\frac{1}{6}n^2+Cn+(\frac{3}{2}-3C)} \\ &\leq q^{\frac{1}{6}n^2+Cn}. \end{aligned}$$

Note that the second inequality is by the induction hypothesis, and the last inequality holds as long as $C \geq 1$.

Now suppose $d = 1$. In this case, Equation 5 is not enough for our purpose. We then need the following refinement. Partition $N(v)$ as $N_1(v) \cup N_{\geq 2}(v)$, where $N_1(v) = \{w \in \mathbb{F}_q^n : w \in N(v), \deg(w) = 1\}$, and $N_{\geq 2}(v) = N(v) \setminus N_1(v)$. A refinement of Equation 4 gives that

$$\text{NMI}(\mathcal{A}) \leq |N_1(v)| \cdot g_q(n-2) + |N_{\geq 2}(v)| \cdot g_q(n-3). \quad (6)$$

If $|N_1(v)| \leq q^{\frac{2}{3}n}$, then we have

$$\begin{aligned} \text{NMI}(\mathcal{A}) &\leq q^{\frac{2}{3}n} \cdot g_q(n-2) + q^n \cdot g_q(n-3) \\ &\leq q^{\frac{2}{3}n} \cdot q^{\frac{1}{6}(n-2)^2 + C(n-2)} + q^n \cdot q^{\frac{1}{6}(n-3)^2 + C(n-3)} \\ &\leq q^{\frac{1}{6}n^2 + Cn + (\frac{2}{3} - 2C)} + q^{\frac{1}{6}n^2 + Cn + (\frac{3}{2} - 3C)} \\ &\leq q^{\frac{1}{6}n^2 + Cn - 1} + q^{\frac{1}{6}n^2 + Cn - 1} \\ &\leq q^{\frac{1}{6}n^2 + Cn}. \end{aligned}$$

Note that the second inequality is by the induction hypothesis, the second to the last inequality holds as long as $C \geq 1$.

If $|N_1(v)| > q^{\frac{2}{3}n}$, then we first prove the following.

Claim 35. *We have $\text{rk}(\mathcal{A}) > \frac{2}{3}n$.*

Proof for Claim 35. Suppose $\dim(\mathcal{A}) = m$. Let $s = \lfloor \frac{2}{3}n \rfloor + 1$, the smallest integer larger than $\frac{2}{3}n$. We will show that there exists a linear basis of some $\tilde{\mathcal{A}}$ that is isometric to \mathcal{A} , $\tilde{A}_1, \dots, \tilde{A}_m \in \Lambda(n, q)$, such that

$$\tilde{A}_1 = \begin{bmatrix} B_1 & -C_1^t \\ C_1 & D_1 \end{bmatrix}, \tilde{A}_2 = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & D_2 \end{bmatrix}, \dots, \tilde{A}_m = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & D_m \end{bmatrix}, \quad (7)$$

where $B_1 \in \Lambda(s, q)$, $C_1 \in M((n-s) \times s, q)$, and $D_i \in \Lambda(n-s, q)$. From this linear basis, it is clear that $\begin{bmatrix} B_1 \\ C_1 \end{bmatrix}$ is of rank $s > \frac{2}{3}n$, as otherwise $\tilde{\mathcal{A}}$ would be degenerate. It would follow then that $\text{rk}(\mathcal{A}) = \text{rk}(\tilde{\mathcal{A}}) > \frac{2}{3}n$.

We first start with an arbitrary linear basis of \mathcal{A} , say $A_1, \dots, A_m \in \Lambda(n, q)$. Recall that v is of degree 1, and $|N_1(v)| > q^{\frac{2}{3}n}$. For later convenience, rename v as u_1 . Then there exist $u_2, u_3, \dots, u_s \in \mathbb{F}_q^n$, such that for $i \geq 2$, $u_i \in N_1(v)$, and u_1, \dots, u_s are linearly independent. As otherwise, suppose the maximum number of linearly independent u_i 's from $N_1(v)$ we can find is $t < s$. Then since $|N_1(v)| > q^{\frac{2}{3}n} \geq q^t = |\langle u_1, \dots, u_t \rangle|$, we can find $u_{t+1} \in N_1(v) \setminus \langle u_1, \dots, u_t \rangle$, a contradiction.

We then can arrange a change of basis matrix T whose first s columns are u_1, \dots, u_s . Apply this change of basis matrix T (by $T^t \cdot T$) to A_1, \dots, A_m to get a linear basis $\bar{A}_1, \dots, \bar{A}_m$ for $\tilde{\mathcal{A}} = T^t \mathcal{A} T$. Recall that e_i denotes the i th standard basis vector. Since u_i 's are of degree 1, for any $i \in [s]$, we have $\tilde{\mathcal{A}}(e_i)$ is of dimension 1. For $2 \leq i \leq s$, since $u_i \in N(v)$, we have

$$e_i^t(\bar{A}_1, \dots, \bar{A}_m)e_1 \neq (0, \dots, 0). \quad (8)$$

Without loss of generality, assume $\bar{A}_1 e_1 \neq \mathbf{0}$. As $\tilde{\mathcal{A}}(e_1)$ is of dimension 1, we have for any $2 \leq j \leq m$, $\bar{A}_j e_1 = \lambda_j \bar{A}_1 e_1$ for some $\lambda_j \in \mathbb{F}_q$. We claim that for any $2 \leq i \leq s$, the i th entry of

$\bar{A}_1 e_1, \bar{A}_1 e_1(i) \neq 0$. If not, then for any $2 \leq j \leq m$, $\bar{A}_j e_1(i) = \lambda_j \bar{A}_1 e_1(i) = 0$. This is equivalent to say that $e_1^t(\bar{A}_1, \dots, \bar{A}_m) e_1 = 0$, contradicting Equation 8.

As \bar{A}_i 's are alternating matrices, we have for any i, j, k , $(\bar{A}_i e_j)(k) = -(\bar{A}_i e_k)(j)$. It follows that for $2 \leq i \leq s$, $\bar{A}_1 e_i(1) = -\bar{A}_1 e_1(i) \neq 0$, and for $2 \leq j \leq m$, $\bar{A}_j e_i(1) = -\bar{A}_j e_1(i) = -\lambda_j \bar{A}_1 e_1(i) = \lambda_j \bar{A}_1 e_i(1)$. Since $\mathcal{A}(e_i)$ is of dimension 1 for $2 \leq i \leq s$, we infer that for $2 \leq j \leq m$ and $2 \leq i \leq s$, $\bar{A}_j e_i = \lambda_j \bar{A}_1 e_i$. We then let $\tilde{A}_1 = \bar{A}_1$, and for $2 \leq j \leq m$, $\tilde{A}_j = \bar{A}_j - \lambda_j \bar{A}_1$. Clearly, $\tilde{A}_1, \dots, \tilde{A}_m$ still form a basis of $\tilde{\mathcal{A}}$, and they are of the form in Equation 7. The claim then follows. \square

Combining Claim 35 and Lemma 34, the proof is concluded. \square

9.2.1 Proof of Lemma 34

Let c be the corank of A . We then have $c < \frac{1}{3}n$.

Let (u_1, \dots, u_d) be an ordered basis of an isotropic space U of A of dimension d . For $i \in [d]$, let $U_i = \langle u_1, \dots, u_i \rangle$, and let $a_i = \dim(A(U_i))$. We also let $U_0 = \mathbf{0}$, and $a_0 = 0$. Note that $U = U_d$, and we also let $a = a_d$. We call such an isotropic space of (d, a) type. Note that $\dim(\langle U_i, \text{rad}(A) \rangle) = \dim(U_i) + \dim(\text{rad}(A)) - \dim(U_i \cap \text{rad}(A)) = \dim(A(U_i)) + \dim(\text{rad}(A)) = a_i + c$.

After fixing u_1, \dots, u_i , a valid u_{i+1} can come from two sources.

1. If $u_{i+1} \notin \langle U_i, \text{rad}(A) \rangle$, then since u_{i+1} needs to satisfy $u_{i+1}^t A u_j = 0$ for $j = 1, \dots, i$, the number of choices of u_{i+1} is upper bounded by $q^{n-a_i} - q^i$.
2. If $u_{i+1} \in \langle U_i, \text{rad}(A) \rangle$, then the number of choices of u_{i+1} is upper bounded by $q^{c+a_i} - q^i$.

So the following indices are important: for $i \in [a]$, let b_i be the smallest $j \in [d]$ such that $a_j = \dim(A(U_j)) = i$. We then have $0 < b_1 < b_2 < \dots < b_a \leq d$. We also let $b_0 = 0$ and $b_{a+1} = d$. We call such an ordered basis of (b_1, \dots, b_a) type of an isotropic space of (d, a) type.

The number of possible types of an isotropic space is trivially upper bounded by n^2 , and the number of possible types of ordered bases of isotropic spaces of type (d, a) is upper bounded by $\binom{d}{a} \leq 2^d \leq 2^n$. So by a multiplicative factor of $n^2 \cdot 2^n$, we can restrict to consider ordered basis (u_1, \dots, u_d) of a fixed type $\mathbf{b} = (b_1, \dots, b_a)$. By the discussion above, if $j = b_i$, then the number of choices for u_j is upper bounded by $t_{d,a,\mathbf{b}}(j) := q^{n-(i-1)} - q^{j-1}$. If $b_i < j < b_{i+1}$, the number of choices for u_j is upper bounded by $t_{d,a,\mathbf{b}}(j) := q^{c+i} - q^{j-1}$. Recall that $(q^n - q^i)/(q^d - q^i) \leq q^{n-d+1}$, for any q and $i < d \leq n$. If $j = b_i$, we have

$$t_{d,a,\mathbf{b}}(j)/(q^d - q^{j-1}) \leq q \cdot q^{n-(i-1)-d}. \quad (9)$$

If $b_i < j < b_{i+1}$, we have

$$t_{d,a,\mathbf{b}}(j)/(q^d - q^{j-1}) \leq q \cdot q^{(c+i)-d} \leq q \cdot q^{(c+a)-d} \quad (10)$$

Each dimension- d subspace of \mathbb{F}_q^n has $(q^d - 1)(q^d - q) \dots (q^d - q^{d-1})$ ordered bases, and each ordered basis of an isotropic space of type (d, a) is of a particular type. The number of dimension- d

isotropic spaces of type (d, a) can be upper bounded by

$$\begin{aligned}
& \sum_{\text{type } \mathbf{b}=(b_1, \dots, b_a)} \frac{t_{d,a,\mathbf{b}}(1) \cdot \dots \cdot t_{d,a,\mathbf{b}}(d)}{(q^d - 1) \cdot \dots \cdot (q^d - q^{d-1})} \\
&= \sum_{\text{type } \mathbf{b}=(b_1, \dots, b_a)} \frac{t_{d,a,\mathbf{b}}(1)}{q^d - 1} \cdot \dots \cdot \frac{t_{d,a,\mathbf{b}}(j)}{q^d - q^{j-1}} \cdot \dots \cdot \frac{t_{d,a,\mathbf{b}}(d)}{q^d - q^{d-1}} \\
&\leq \sum_{\text{type } \mathbf{b}=(b_1, \dots, b_a)} q^d \cdot q^{na - \sum_{i=1}^a (i-1) - da} \cdot q^{(c+a)(d-a) - d(d-a)} \\
&\leq 2^n \cdot q^{na - a^2/2 + (c+a)(d-a) - d^2 + d + a/2}.
\end{aligned}$$

Let us explain the first inequality. The q^d term is because of the q terms on the right hand sides of Equations 9 and 10. The $q^{na - \sum_{i=1}^a (i-1) - da}$ is by collecting those terms from Equation 9, and the $q^{(c+a)(d-a) - d(d-a)}$ term is by collecting those terms from Equation 10.

It is then clear that we need to bound $f(n, d, a) = na - a^2/2 + (c+a)(d-a) - d^2$ for $1 \leq a \leq d \leq n$. After some arrangement, we have

$$f(n, d, a) = -\frac{3}{2}\left(a - \frac{1}{3}(n + d - c)\right)^2 + \frac{1}{6}(n + d - c)^2 - d^2 + dc.$$

We then distinguish between two cases.

1. Case (i): when $\frac{1}{3}(n + d - c) \leq d$ holds, namely $d \geq (n - c)/2$. Only in this case, a can be set to $\frac{1}{3}(n + d - c)$, and the maximum can be set to $g(n, d) := \frac{1}{6}(n + d - c)^2 - d^2 + dc$. After some arrangement, we have

$$g(n, d) = -\frac{5}{6}\left(d - \frac{1}{5}(n + 2c)\right)^2 + \frac{1}{30}(n + 2c)^2 + \frac{1}{6}(n - c)^2.$$

Since $c < n/3$, we have $(n - c)/2 > (n + 2c)/5$. Recall that $d \geq (n - c)/2$. So $g(n, d)$ achieves maximal at $d = (n - c)/2$. Plugging this in, the maximal value is

$$h(n) := g(n, (n - c)/2) = -\frac{3}{8}\left(c - \frac{1}{3}n\right)^2 + \frac{1}{6}n^2 < \frac{1}{6}n^2.$$

2. Case (ii): when $\frac{1}{3}(n + d - c) > d$ holds, namely $d < (n - c)/2$. In this case, $f(n, d, a)$ achieves the maximal value at $a = d$, and

$$f(n, d, d) = -\frac{3}{2}\left(d - \frac{1}{3}n\right)^2 + \frac{1}{6}n^2 \leq \frac{1}{6}n^2,$$

where the inequality becomes an equality at $d = n/3$.

Since in both cases, the maximal value is no more than $\frac{1}{6}n^2$, we can then conclude the proof.

9.3 Turning Theorem 11 into an algorithm

The proof of Theorem 11 can be turned into an algorithm for enumerating all maximal isotropic spaces in time $q^{\frac{1}{6}n^2 + O(n)}$. We briefly indicate some algorithmic issues for doing this. Firstly, note that in time $q^{O(n)}$, one can compute $\deg_{\mathcal{A}}(v)$ for all $v \in \mathbb{F}_q^n$. Secondly, the Equation 4 naturally

suggests a recursive algorithm structure. In the cases when $d \geq 2$, or $d = 1$ and $|N_1(v)| \leq q^{\frac{2}{3}n}$, this recursive structure readily gives the desired algorithm. If $d = 1$ and $|N_1(v)| > q^{\frac{2}{3}n}$, we need to make the proofs of Claim 35 and Lemma 34 constructive. Then for each isotropic space of some $A \in \Lambda(n, \mathbb{F})$, A of rank $> \frac{2}{3}n$, test whether it is maximal using Observation 19.

For Claim 35, note that the selection of u_i 's from $N_1(v)$ can be done easily in a greedy way. Other steps are readily constructive. For Lemma 34, we use the same procedure as in Proposition 32, whose running time is bounded in time $q^{\frac{1}{6}n^2 + O(n)}$ by Lemma 34.

We then have the following.

Corollary 36. *Given $\mathcal{A} \leq \Lambda(n, q)$, all maximal isotropic spaces can be enumerated in time $q^{\frac{1}{6}n^2 + O(n)}$.*

10 Proof of Theorem 12

Review of Lawler's algorithm. We first review Lawler's dynamic programming idea for computing the chromatic number [Law76], and then adapt that idea to our problem.

Given a graph $G = (V, E)$, Lawler's algorithm for computing $\chi(G)$ goes as follows. The idea is to build a table storing $\chi(H)$ for every induced subgraph H of G . Note that this table is of size 2^n . To fill in this table, the starting point is the empty graph with chromatic number 0. Suppose we have computed the chromatic numbers of those induced subgraphs of size $< \ell$. Let $H = (U, F)$ be an induced subgraph of size ℓ . Then the chromatic number of H can be computed by the following formula:

$$\chi(H) = 1 + \min_{I \subseteq U} \{\chi(H[U \setminus I])\},$$

where I goes over all maximal independent sets of H , and $H[U \setminus I]$ is the induced subgraph of H restricting to vertex set $U \setminus I$. Since there are at most $3^{\ell/3}$ maximal independent sets of H and they can be enumerated in time $O(3^{\ell/3} \cdot n)$, the exponential part of the time complexity of this algorithm is $\sum_{\ell=0}^n \binom{n}{\ell} \cdot 3^{\ell/3} = (1 + \sqrt[3]{3})^n$.

Directly applying Lawler's idea to isotropic numbers. The above idea can be adapted to compute $\chi(\mathcal{A})$ for $\mathcal{A} \leq \Lambda(n, q)$ as follows. To start with, recall that in the above algorithm we used the following simple fact: if a graph G admits a vertex c -coloring, then there is a vertex c -coloring in which one part is a maximal independent set. We leave the reader to check that the analogue of this fact in the alternating matrix space setting also holds.

Given $\mathcal{A} \leq \Lambda(n, q)$, we also store a table storing $\chi(\mathcal{B})$ for every induced alternating matrix spaces \mathcal{B} of \mathcal{A} . Note that this table is of size $q^{\frac{1}{4}n^2 + O(n)}$. To fill in this table, the starting point is the zero space with isotropic decomposition number 0. Suppose we have computed the isotropic decomposition numbers of those induced alternating matrix spaces of dimension $< \ell$. Let $\mathcal{B} \leq \Lambda(\ell, q)$ be an induced alternating matrix space corresponding to $U \leq \mathbb{F}_q^n$ of dimension ℓ . Then the isotropic decomposition number of \mathcal{B} can be computed by the following formula

$$\chi(\mathcal{B}) = 1 + \min_{V \subseteq U, W \subseteq U} \{\chi(\mathcal{B}|_W)\},$$

where V goes over all maximal isotropic spaces of \mathcal{B} , and W goes over all complement subspaces of V in U . Note that here we also need to enumerate all complements of V , while in the graph setting,

the complement set is unique. Recall that by Theorem 11, there are at most $q^{\frac{1}{6}\ell^2+O(\ell)}$ maximal isotropic spaces of \mathcal{B} , and they can be enumerated in time $q^{\frac{1}{6}\ell^2+O(\ell)}$. This gives a bound on the number of V . We bound the number of W using the trivial $q^{\frac{1}{4}\ell^2+O(\ell)}$ bound. Note that we will need to test whether W is a complement of V , which can be done easily. Another more efficient approach would be to enumerate all complements of U in time $q^{d(n-d)} \cdot \text{poly}(n, \log q)$ (see [LQ17, Proposition 17 in the arXiv version]).

So to fill in those entries corresponding to alternating matrix spaces induced by ℓ -dimensional subspaces, the time complexity is bounded by

$$\begin{aligned}
& \binom{n}{\ell}_q \cdot q^{\frac{1}{6}\ell^2+O(\ell)} \cdot q^{\frac{1}{4}\ell^2+O(\ell)} \\
& \leq q^{\ell(n-\ell)+\ell} \cdot q^{\frac{1}{6}\ell^2+O(\ell)} \cdot q^{\frac{1}{4}\ell^2+O(\ell)} \\
& = q^{\ell n - \frac{7}{12}\ell^2 + O(\ell)} \\
& = q^{-\frac{7}{12}(\ell - \frac{6}{7}n)^2 + \frac{3}{7}n^2 + O(\ell)} \\
& \leq q^{\frac{3}{7}n^2 + O(n)}.
\end{aligned}$$

Summing over $\ell \in \{0, 1, \dots, n\}$, we see that the total time complexity is also bounded by $q^{\frac{3}{7}n^2 + O(n)}$.

A new dynamic programming scheme. In the above, we see that directly following Lawler's dynamic programming scheme does lead to an improved algorithm for computing the isotropic decomposition number. However, a key difference with the classical setting, namely the magnitude of complement subspaces, impacts the analysis. In the following, we shall use another dynamic programming scheme, still combined with the $q^{\frac{1}{6}n^2+O(n)}$ upper bound on the number of maximal isotropic spaces, to achieve the $q^{\frac{5}{12}n^2+O(n)}$ running time as promised in Theorem 12.

To do that, we first make a simple observation.

Observation 37. *Let $\mathcal{A} \leq \Lambda(n, \mathbb{F})$. Then $\chi(\mathcal{A}) \leq k$, if and only if, there exist k maximal isotropic spaces U_1, \dots, U_k , such that $\mathbb{F}^n = \langle \cup_{i \in [k]} U_i \rangle$.*

Proof. For the only if direction, recall that every isotropic space is contained in a maximal one. For the if direction, note that from U_1, \dots, U_k , we can construct U'_1, \dots, U'_k , such that $U'_i \leq U_i$, and U'_1, \dots, U'_k form a direct sum decomposition of \mathbb{F}^n . This shows that $\chi(\mathcal{A}) \leq k$. \square

The key to our algorithm is the following function. For $k \in [n]$ and $W \leq \mathbb{F}_q^n$, let $f(k, W)$ be the boolean function such that $f(k, W) = 1$ if and only if $W = \langle \cup_{i \in [k]} U_i : U_i \text{ maximal isotropic} \rangle$. For example, $f(1, W) = 1$ if and only if W is a maximal isotropic space.

The following is then a dynamic programming scheme computing $f(k, W)$ for every $k \in [n]$ and $W \leq \mathbb{F}_q^n$. Let $\mathcal{A} \leq \Lambda(n, q)$ be an alternating matrix space. We assume that $\chi(\mathcal{A}) > 1$, as $\chi(\mathcal{A}) = 1$ if and only if \mathcal{A} is the zero space.

1. Use Corollary 36 to compute the set of maximal isotropic spaces of \mathcal{A} , and let MI be this set.
2. Build a table f , indexed by (k, W) where $k \in [n]$ and $W \leq \mathbb{F}_q^n$, and initiate $f(k, W) = 0$ for every k and W .
3. For every $W \leq \mathbb{F}_q^n$, do:

(a) If W is maximal isotropic, then $f(1, W) = 1$.

4. For $k = 2, \dots, n$, do:

(a) For every $W \leq \mathbb{F}_q^n$ and every $T \in MI$, do:

i. If $f(k-1, W) = 1$, then let $U = \langle W \cup T \rangle$, and set $f(k, U) = 1$.

ii. If $U = \mathbb{F}_q^n$, then return “ $\chi(\mathcal{A}) = k$.”

To prove the correctness of the algorithm, we first note that by induction, the algorithm correctly computes $f(k, W)$ for every k and W . Then suppose the algorithm returns with reporting that $\chi(\mathcal{A}) = k$. Note that in this case, it does find k maximal isotropic subspaces covering the whole space \mathbb{F}_q^n . So by Observation 37, $\chi(\mathcal{A}) \leq k$. So we are left to show that $\chi(\mathcal{A}) \geq k$. By way of contradiction, suppose $\chi(\mathcal{A}) = k' < k$, so by Observation 37, there exist maximal isotropic subspaces $U_1, \dots, U_{k'}$ that cover \mathbb{F}_q^n . Let $W = \langle U_2 \cup \dots \cup U_{k'} \rangle$. Then W is a proper subspace of \mathbb{F}_q^n , as otherwise by Observation 37 $\chi(\mathcal{A}) \leq k' - 1 < k'$, contradicting that $\chi(\mathcal{A}) = k'$. But this means that $f(k' - 1, W) = 1$, so in Step (4.a), when enumerating W and $T = U_1$, the algorithm would go through steps (4.a.i) and (4.a.ii), and outputs that $\chi(\mathcal{A}) = k'$. This gives us the desired contradiction.

To estimate the running time of the algorithm, note that Step (1) costs $q^{\frac{1}{6}n^2 + O(n)}$ by Corollary 36. All subspaces can be enumerated in time $q^{\frac{1}{4}n^2 + O(n)}$ by the same technique as in the proof of Proposition 32. The total running time is then dominated by the loop in steps (4) and (4.a), which is $n \cdot q^{\frac{1}{6}n^2 + O(n)} \cdot q^{\frac{1}{4}n^2 + O(n)} = q^{\frac{5}{12}n^2 + O(n)}$. This concludes the proof of Theorem 12.

11 Proofs for Propositions 13 and 14

Proof of Proposition 13. When $\mathbb{F} = \mathbb{C}$ and the input instance is over \mathbb{Z} , we shall formulate the maximum isotropic space problem as a problem about the solvability of a system of integral polynomial equations over \mathbb{C} . The result would follow then by using Koiran’s result that the Hilbert Nullstellensatz is in PH, assuming the generalized Riemann hypothesis [Koi96]. We first cite Koiran’s result as follows, following the formulation of [Mul17, Theorem 2.10].

Theorem 38 ([Koi96]). *The problem Hilbert’s Nullstellensatz of deciding whether a given system of multivariate integral polynomials, specified as arithmetic circuits, has a solution over \mathbb{C} is in PSPACE unconditionally, and in $RP^{NP} \subseteq \Pi_2$ assuming the generalized Riemann hypothesis.*

Therefore, to put the maximum isotropic space problem over \mathbb{C} in PSPACE unconditionally, and in PH assuming the generalized Riemann hypothesis, for instances given by integral alternating matrices, we only need to formulate this problem as deciding the solvability of a system of integral polynomial equations represented by arithmetic circuits. This can be done as follows. Suppose we are given $\mathcal{A} = \langle A_1, \dots, A_m \rangle \leq \Lambda(n, \mathbb{C})$ where A_i ’s are integral matrices, and we want to know whether there exists an isotropic space of dimension d for \mathcal{A} . Then \mathcal{A} has an isotropic space of dimension d if and only if there exists an invertible matrix T such that for any $i \in [m]$, the left-upper $d \times d$ submatrix of $T^t A_i T$ consists of all zero entries. We then set up an $n \times n$ variable matrix $X = (x_{i,j})_{i,j \in [n]}$, and a variable y . The role of this variable y is to impose the condition that $\det(X)$ is non-zero, as will be seen below soon. For every $i \in [m]$, set the entries of the left-upper $d \times d$ submatrix of $X^t A_i X$ to be zero. This gives md^2 integral quadratic polynomials in $x_{i,j}$ ’s. To enforce that the valid solutions are from invertible matrices, we set up the equation $\det(X) \cdot y = 1$,

which is also an integral polynomial. Note that the polynomial $\det(X)$ can be expressed as a small arithmetic circuit. It is straightforward to verify that these $(md^2 + 1)$ equations in $x_{i,j}$ and y have a non-trivial solution if and only if \mathcal{A} has a dimension- d isotropic space.

Let c be a constant. For isotropic c -decomposition problem, the idea is basically the same. The only small complication is that we need to specify the dimensions of the three isotropic spaces in a c -isotropic decomposition. But the number of possibilities is at most n^c , which we can enumerate as c is a constant. After fixing some (d_1, d_2, \dots, d_c) , where $d_i \in \mathbb{Z}^+$, $n \geq d_1 \geq \dots \geq d_c \geq 1$, and $\sum_{i \in [c]} d_i = n$, we can construct a system of integral polynomial equations to express the condition that there exists a c -isotropic decomposition with these dimensions, just as in the case of the maximum isotropic space problem. This concludes the proof. \square

Proof of Proposition 14. Recall that we have $\mathcal{A} \leq \Lambda(n, \mathbb{F})$, and our goal is to prove $\chi(\mathcal{A}) \leq O(\Delta(\mathcal{A}) \cdot \log n)$. Here, $\Delta(\mathcal{A}) := \max\{\deg_{\mathcal{A}}(v) : v \in \mathbb{F}^n\}$, and $\deg_{\mathcal{A}}(v) := \dim(\langle Av : A \in \mathcal{A} \rangle)$. We will also use a greedy algorithm to construct an isotropic C -decomposition with $C \leq O(\Delta(\mathcal{A}) \cdot \log n)$.

For $v \in \mathbb{F}^n$, recall that $\text{rad}_{\mathcal{A}}(v) = \{u \in \mathbb{F}^n : u^t Av = 0\}$. Consider the following algorithm.

1. Set $k = 0$, and $U = \mathbf{0}$;
2. While $\dim(U) < n$, do:
 - (a) Set $k = k + 1$;
 - (b) Let W be any complementary subspace of U ;
 - (c) Let $S = \emptyset$;
 - (d) While $\dim(W) > |S|$, do:
 - i. Take any $w \in W \setminus \langle S \rangle$; // $\langle \emptyset \rangle := \mathbf{0}$
 - ii. $S = S \cup w$;
 - iii. $W = W \cap \text{rad}_{\mathcal{A}}(w)$;
 - (e) Let $U_k = \langle S \rangle$;
 - (f) $U = \langle U \cup U_k \rangle$;
3. Return $U_1 \oplus U_2 \oplus \dots \oplus U_k$.

We first argue that $U_1 \oplus U_2 \oplus \dots \oplus U_k$ is an isotropic k -decomposition of \mathcal{A} . To see this, we note that because of Step (2.d.iii), the condition $W \subseteq \text{rad}(\langle S \rangle)$ holds in the loop of Step (d), so $\langle S \rangle$ maintains as an isotropic space by Observation 19.

We then show that $k = O(\Delta(\mathcal{A}) \cdot \log n)$ when the algorithm terminates. For this, let $d_i = \dim(U_i)$, and $D_i = d_1 + \dots + d_i$. Observe that $\dim(\text{rad}_{\mathcal{A}}(w)) \geq n - \Delta(\mathcal{A})$. It follows that $\dim(W \cap \text{rad}_{\mathcal{A}}(w)) = \dim(W) + \dim(\text{rad}_{\mathcal{A}}(w)) - \dim(\langle W \cup \text{rad}_{\mathcal{A}}(w) \rangle) \geq \dim(W) + (n - \Delta(\mathcal{A})) - \dim(\langle W \cup \text{rad}_{\mathcal{A}}(w) \rangle) \geq \dim(W) - \Delta(\mathcal{A})$. Therefore, in the computing procedure of U_i , we have $d_i = \dim(U_i) \geq \frac{n - D_{i-1}}{\Delta(\mathcal{A})}$. This implies that $n - D_i = n - D_{i-1} - d_i \leq (n - D_{i-1})(1 - 1/\Delta(\mathcal{A}))$. Therefore, adding a new U_i to the direct sum decomposition decreases the value of $n - D_i$ by a factor of at least $1 - 1/\Delta(\mathcal{A})$. Therefore the algorithm terminates in at most $\log_{1-1/\Delta(\mathcal{A})} \frac{1}{n} = O(\Delta(\mathcal{A}) \cdot \log n)$ steps. \square

12 Proofs of Theorems 15 and 16

In this section we prove theorems 15 and 16. While the proofs are straightforward for experts, we include details for completeness. We shall refer to some facts in Section 4.1 from time to time.

Proof of Theorem 15. Recall that the goal is to show that deciding whether a matrix group contains an abelian subgroup of order $\geq s$ is NP-hard for some $s \in \mathbb{N}$. We shall reduce the maximum isotropic space problem, which is NP-hard by Corollary 4, to this problem.

For this we shall need the following classical construction¹¹. Let p be an odd prime, and let $\mathcal{A} \leq \Lambda(n, p)$ be given by an ordered linear basis $\mathbf{A} = (A_1, \dots, A_m)$. Recall that e_i denotes the i th standard basis vector. From \mathbf{A} , for $i \in [n]$, construct $B_i = [A_1 e_i, \dots, A_m e_i] \in M(n \times m, p)$. That is, the j th column of B_i is the i th column of A_j . Then for $i \in [n]$, construct

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

and for $j \in [m]$, construct

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

Let $G_{\mathbf{A}}$ be the matrix group generated by \tilde{B}_i and \tilde{C}_j . Then it can be verified easily that, $G_{\mathbf{A}}$ is isomorphic to the Baer group (see Section 4.1) corresponding to the alternating bilinear map defined by \mathbf{A} (see Remark 17). In particular, $[G, G] \cong \mathbb{Z}_p^m$, and $G/[G, G] \cong \mathbb{Z}_p^n$. By the correspondence between isotropic spaces of \mathcal{A} and abelian normal subgroups of $G_{\mathbf{A}}$ containing the commutator subgroup (see Section 4.1), deciding whether \mathcal{A} has an isotropic space of dimension $\geq d$ is equivalent to deciding whether $G_{\mathbf{A}}$ has an abelian subgroup of order $\geq s = p^{m+d}$. This completes the reduction. \square

Proof of Theorem 16. Let P be a p -group of class 2 and exponent p , and let $\phi : P/[P, P] \times P/[P, P] \rightarrow [P, P]$ be the commutator map. The proof of Theorem 16 basically follows from the correspondence between abelian subgroup containing $[P, P]$ and isotropic spaces of ϕ as described in Section 4.1. The only small caveat here is that we need a bound on the dimension of $P/Z(P)$ instead of the dimension of $P/[P, P]$. To overcome this, we first observe that a maximal abelian subgroup of P necessarily contains the center $Z(P)$, which in turn contains $[P, P]$ by the class-2 condition. Then we only need to note that $Z(P)/[P, P]$ corresponds to the radical of ϕ , and recall that the number of maximal isotropic spaces only depends on the non-degenerate part of ϕ by Observation 33. The proof then can be concluded. \square

13 A quantum variant of the theory

One way to extend isotropic spaces and isotropic decompositions to the quantum information setting is as follows. Briefly speaking, firstly we restrict to the complex number field \mathbb{C} . Secondly, instead

¹¹We thank James B. Wilson for communicating this construction to us.

of tuples of alternating matrices, we will consider tuples of matrices which represent a so-called irreducible quantum channels. Thirdly, instead of general linear groups, we will consider unitary groups.

For detailed explanations, we need some notation. For $A \in M(n, \mathbb{C})$ we use $A \succeq 0$ to denote that A is positive semi-definite, and $A \succ 0$ to denote that A is positive definite. For $\mathbf{B} = \{B_1, \dots, B_m\} \subseteq M(n, \mathbb{C})$, we let $\tilde{\mathbf{B}} : M(n, \mathbb{C}) \rightarrow M(n, \mathbb{C})$ be the function sending $A \in M(n, \mathbb{C})$ to $\sum_{i=1}^m B_i A B_i^\dagger$. It is clear that $\tilde{\mathbf{B}}$ can be represented as an $n^2 \times n^2$ matrix $\sum_{i=1}^m B_i \otimes B_i^*$, where B_i^* stands for the entry-wise complex conjugation of B_i .

Let $D(n, \mathbb{C}) \subseteq M(n, \mathbb{C})$ be the set of $n \times n$ positive semi-definite matrices with unit trace over \mathbb{C} , and let $D^+(n, \mathbb{C}) \subseteq D(n, \mathbb{C})$ consist of those positive definite matrices in $D(n, \mathbb{C})$. Elements from $D(n, \mathbb{C})$ are known as quantum states.

Let $\text{QC}(n, \mathbb{C})$ be the set of sets of matrices $\mathbf{B} = \{B_1, \dots, B_m\} \subseteq M(n, \mathbb{C})$ satisfying $\sum_{i=1}^m B_i^\dagger B_i = I$. Functions of the form $\tilde{\mathbf{B}}$ for $\mathbf{B} \in \text{QC}(n, \mathbb{C})$ are known as quantum channels, as they are completely positive and trace preserving.

We then define isotropic spaces and decompositions in the quantum setting. To define isotropic spaces, we essentially follow the same pattern as in the alternating matrix space setting. For isotropic decompositions, we shall require that the direct sum decomposition is also an orthogonal one, as the underlying spaces of quantum channels are Hilbert spaces which come with a norm.

Definition 39. Let $\mathbf{B} = \{B_1, \dots, B_m\} \in \text{QC}(n, \mathbb{C})$. An isotropic space of \mathbf{B} is a subspace $U \leq \mathbb{C}^n$, such that for any $u, u' \in U$, and any B_i , we have $u^\dagger B_i u' = 0$. An isotropic c -decomposition of \mathbf{B} is an orthogonal direct sum decomposition of $\mathbb{C}^n = U_1 \oplus U_2 \oplus \dots \oplus U_c$ such that each U_i is a non-zero isotropic space of \mathbf{B} .

13.1 From connected graphs to irreducible quantum channels

In this subsection, we establish a connection between independent sets and vertex colorings of connected graphs, and isotropic spaces and decompositions of a particular type of quantum channels, called irreducible channels (defined below).

We obtain two main results. The first result, Proposition 41, reduces certain problems for connected graphs to the corresponding ones for irreducible quantum channels. This result corresponds to Theorem 3. The second result, Theorem 43, gives an efficient algorithm for isotropic 2-decomposition in this setting. This result corresponds to Theorem 6, but the techniques are completely different.

Let $\text{IQC}(n, \mathbb{C}) \subseteq \text{QC}(n, \mathbb{C})$ consist of those $\mathbf{B} \in \text{QC}(n, \mathbb{C})$ satisfying the following: there exists a unique fixed $\rho \in D(n, \mathbb{C})$ of $\tilde{\mathbf{B}}$, and further $\rho \in D^+(n, \mathbb{C})$, where ρ is said to be fixed of $\tilde{\mathbf{B}}$ if $\tilde{\mathbf{B}}(\rho) = \rho$. Such \mathbf{B} and $\tilde{\mathbf{B}}$ are called *irreducible*. Irreducible quantum channels have been studied in e.g. [Dav70] and [Wol12, Sec. 6.2]. In particular, the definition of irreducible quantum channels follows from [Dav70, Theorem 13]. Furthermore, given $\mathbf{B} \in \text{QC}(n, \mathbb{C})$, let M be the $n^2 \times n^2$ matrix representation of $\tilde{\mathbf{B}}$. Then $\mathbf{B} \in \text{IQC}(n, \mathbb{C})$ if and only if the algebraic and geometric multiplicities of the eigenvalue 1 of M are both 1, and any 1-eigenvector is of full-rank.

We first observe that a simple and connected graph can be realized as an irreducible quantum channel as follows. This is classical, but for completeness we spell out the details. Let $G = ([n], E)$ be a connected graph. For each $i \in [n]$, let d_i be the degree of i . We construct the following set of matrices $\mathbf{B}_G = \{\frac{1}{\sqrt{d_j}} \cdot E_{i,j}, \frac{1}{\sqrt{d_i}} \cdot E_{j,i} : \{i, j\} \in E\}$. Note that $|\mathbf{B}_G| = 2|E|$.

Proposition 40. *Let G and \mathbf{B}_G be as above. Then $\mathbf{B}_G \in \text{IQC}(n, \mathbb{C})$.*

Proof. We first verify that $\tilde{\mathbf{B}}_G$ is a quantum channel. For this, observe that $(\frac{1}{\sqrt{d_j}} \cdot E_{i,j})^\dagger \frac{1}{\sqrt{d_j}} \cdot E_{i,j} = \frac{1}{d_j} E_{j,i} E_{i,j} = \frac{1}{d_j} E_{j,j}$. Since each vertex i connecting to j contributes one such term, and it follows that $\sum_{E \in \mathbf{B}_G} E^\dagger E = I$. We then verify that $\tilde{\mathbf{B}}_G$ is irreducible. For this, consider $P = (p_{i,j})$ where $p_{i,j} = 1/d_i$, which represents the transition matrix of the Markov chain naturally associated with G . Since G is connected, this Markov chain is irreducible, so there exists a unique probability distribution, e.g. a row vector $s = (s_1, \dots, s_n)$ satisfying $s_i > 0$, $\sum_i s_i = 1$, such that $sP = s$ (see e.g. [LP17, Corollary 1.17]). It can then be verified that the matrix $S = \text{diag}(s_1, \dots, s_n) \in D^+(n, \mathbb{C})$ is fixed by $\tilde{\mathbf{B}}_G$. To see that this is the unique fixed state, we represent $\tilde{\mathbf{B}}_G$ as an $n^2 \times n^2$ matrix M_G . It is not hard to see that by conjugating with a permutation matrix, M_G is of the form $\begin{bmatrix} P & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$. Therefore, the algebraic and geometric multiplicities of the eigenvalue 1 of M_G are the same as those for P , which are 1 by the Perron-Frobenius theory. It follows that \mathbf{B} is irreducible. \square

Proposition 41. *Let $G = ([n], E)$ be a connected graph, and let $\mathbf{B}_G \in \text{IQC}(n, \mathbb{C})$ be as above.*

1. *G has a size- s independent set if and only if \mathbf{B}_G has a dimension- s isotropic space;*
2. *G has a vertex c -coloring if and only if \mathbf{B}_G has an isotropic c -decomposition.*

Proof. (1) The only if direction is trivial. For the if direction, let U be a dimension- s isotropic space of \mathbf{B}_G . Then U is also a dimension- s isotropic space of the alternating matrix space $\langle E_{i,j} - E_{j,i} : \{i, j\} \in E \rangle$, because it is a subspace of $\langle \mathbf{B}_G \rangle$. We can then conclude by resorting to Theorem 3.

(2) The only if direction is trivial; observe that the direct sum decomposition obtained from a vertex coloring as in Theorem 3 is also an orthogonal direct sum decomposition. For the if direction, we observe that, an orthogonal direct sum decomposition into isotropic spaces for \mathbf{B}_G is also one for the alternating matrix space $\langle E_{i,j} - E_{j,i} : \{i, j\} \in E \rangle$. We can then conclude by resorting to Theorem 3. \square

Since the maximum independent set problem and the vertex 3-coloring problem on connected graphs are also NP-hard, we have the following.

Corollary 42. *The maximum isotropic space problem and the isotropic 3-decomposition problem for $\mathbf{B} \in \text{IQC}(n, \mathbb{C})$ are NP-hard.*

This also leaves the isotropic 2-decomposition problem an interesting question. For this, we can resort to the techniques developed for quantum Markov chains, mostly notably, based on recent works of periodicity of quantum channels [GFY18].

Theorem 43. *Suppose we are given $\mathbf{B} \in \text{IQC}(n, \mathbb{C})$ such that every matrix in \mathbf{B} are over \mathbb{Q} . There exists an algorithm that decides whether \mathbf{B} admits an isotropic 2-decomposition in polynomial time.*

Proof. The key observation is to characterize isotropic 2-decompositions using the *periodicity* of irreducible quantum channels.

Definition 44 ([FP09]). *Given $\mathbf{B} \in \text{IQC}(n, \mathbb{C})$, the period of \mathbf{B} is the maximum integer m for which there exists an orthogonal direct sum decomposition $\mathbb{C}^n = U_1 \oplus \dots \oplus U_m$ such that for any $i \in [m]$, and any $B \in \mathbf{B}$, we have $B(U_{i \boxminus 1}) \leq U_i$, where \boxminus indicates subtraction modulo m in the range of $[m]$.*

The following lemma relates isotropic 2-decompositions with periodicity.

Lemma 45. *Given $\mathbf{B} \in \text{IQC}(n, \mathbb{C})$, \mathbf{B} admits an isotropic 2-decomposition if and only if the period of \mathbf{B} is $2k$ for some integer k .*

Proof. For the if direction, let $\mathbb{C}^n = U_1 \oplus U_2 \oplus \dots \oplus U_{2k}$ be the orthogonal direct sum decomposition corresponding to the period of \mathbf{B} . Let $V_1 = \langle U_i : i = 2j - 1, j \in [k] \rangle$, and $V_2 = \langle U_i : i = 2j, j \in [k] \rangle$. Then $V_1 \oplus V_2$ is an orthogonal direct sum decomposition, and for any $B \in \mathbf{B}$, $B(V_1) \leq V_2$, and $B(V_2) \leq V_1$. By the orthogonal condition, $v_1^\dagger v_2 = 0$ for any $v_1 \in V_1, v_2 \in V_2$. We then have for any $i = 1, 2$, any $v_i, v'_i \in V_i$, and any $B \in \mathbf{B}$, we have $v_i^\dagger B v'_i = 0$. That is, V_1 and V_2 are isotropic spaces. It follows that the direct sum decomposition $V_1 \oplus V_2$ is an isotropic 2-decompositions.

For the only if direction, let $\mathbb{C}^n = V_1 \oplus V_2$ be an isotropic 2-decomposition. Let P_1 be the projection into V_1 along V_2 , and P_2 the projection into V_2 along V_1 . We have $P_1 + P_2 = I$, and $P_i^\dagger = P_i$. Since V_1 and V_2 are isotropic spaces, for any $B \in \mathbf{B}$, and any $i = 1, 2$, $P_i B P_i = \mathbf{0}$. Using $P_1 + P_2 = I$, it follows that $P_2 B = B P_1$, and $P_1 B = B P_2$. We are then in the position to apply [FP09, Lemma 4.2], to conclude that the period of \mathbf{B} is $2k$ for some integer k . \square

Given Lemma 45, it is enough to compute the period of \mathbf{B} , and this can be done by resorting to the algorithm in [GFY18]. For completeness, we give a brief sketch of the idea. By Lemma 13 of [GFY18], the period of irreducible quantum channel is equivalent to be the number of eigenvalues with magnitude one of the quantum channel. Using the terminologies in the present article, we have the following lemma.

Lemma 46 ([GFY18, Lemma 13]). *Given $\mathbf{B} \in \text{IQC}(n, \mathbb{C})$, the period of \mathbf{B} is equal to the number of eigenvalues of $\tilde{\mathbf{B}}$ with magnitude one.*

Given this lemma, we can explicitly write out the form of $\tilde{\mathbf{B}}$ as an $n^2 \times n^2$ matrix, and compute its eigenvalues using e.g. [Cai94] in the exact model (Section 3.1). Therefore, each eigenvalue α is represented by an irreducible polynomial $f(x)$ and a separating rectangle in the complex plane. To decide whether α has magnitude 1 can be done efficiently by resorting to techniques from [Lov86]. \square

Finally, we remark that the investigation in this subsection is not completely satisfactory. It would be more satisfying to consider isotropic spaces and isotropic decompositions for arbitrary quantum channels, not just the irreducible ones. We adopt the current strategy, partly because for irreducible channels, the periodicity is well-studied and well-connected with isotropic 2-decomposition. We leave it a future work to study isotropic spaces and decompositions in the general setting.

13.2 Quantum gate subspace-fidelity and isotropic spaces

We provide one quantum information theoretic interpretation for isotropic spaces, by relating it to quantum gate (state) fidelity [NC02, Section 9] and noiseless subspaces in quantum error correction [KL97, Lid14]. For the sake of readers who have little quantum information knowledge, we shall proceed by introducing all the necessary notions from quantum information, even though most of them are standard.

In quantum information theory, the fidelity is a measure of the “closeness” of two quantum states, generalizing the fidelity of two distributions over finite events. It expresses the probability

that one state will pass a test (quantum measurements) to identify as the other. Formally, the fidelity of two quantum states $\rho, \sigma \in \mathcal{D}(n, \mathbb{C})$ is defined by

$$F(\rho, \sigma) = [\text{tr}(\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}})]^2.$$

It is worth noting that $0 \leq F(\rho, \sigma) \leq 1$. Furthermore,

- $F(\rho, \sigma) = 0$ if and only if ρ and σ are orthogonal, i.e., $\text{tr}(\rho\sigma) = 0$;
- $F(\rho, \sigma) = 1$ if and only if $\rho = \sigma$.

Quantum state fidelity induces quantum gate fidelity. Unitary channels (i.e. channels of the form $\tilde{V}(A) = VAV^\dagger$ for some unitary matrix $V \in M(n, \mathbb{C})$) are exactly the channels that do not introduce mixedness (i.e., decoherence) into states. Therefore, in experimental settings, they are considered to be the ideal type of channels to be implemented [NC02, Section 8]. However, no implementation of a channel is perfect, as there is no closed (isolated) system, so environment errors are unavoidable, which cause the channel actually implemented to be not unitary. The gate fidelity is a tool for comparing how well the implemented quantum channel $\tilde{\mathbf{B}}$ approximates the desired unitary channel \tilde{V} . Specifically, the gate fidelity on a pure state (uu^\dagger for a normalized vector $u \in \mathbb{C}^n$) is a function defined as follows:

$$F_{\tilde{\mathbf{B}}, \tilde{V}}(u) = F(\tilde{\mathbf{B}}(uu^\dagger), Vuu^\dagger V^\dagger) = u^\dagger V^\dagger \tilde{\mathbf{B}}(uu^\dagger) V u = u^\dagger [\tilde{V}^\dagger \circ \tilde{\mathbf{B}}(uu^\dagger)] u,$$

where $\tilde{V}^\dagger(A) = V^\dagger A V$. In particular, $F_{\tilde{\mathbf{B}}, \tilde{V}}(u) = F_{\tilde{V}^\dagger \circ \tilde{\mathbf{B}}, \tilde{I}}(u)$, where \tilde{I} is the identity channel. Then the gate fidelity on all states is defined as follows:

$$F(\tilde{\mathbf{B}}) = \min_{u \in \mathbb{C}^n} F_{\tilde{\mathbf{B}}, \tilde{I}}(u) = \min_{\rho \in \mathcal{D}(n, \mathbb{C})} F_{\tilde{\mathbf{B}}, \tilde{I}}(\rho)$$

The second equation in the above follows from the joint concavity property of the state fidelity F (see [NC02, Equation 9.121]).

As we can see, quantum gate fidelity is a global property over \mathbb{C}^n . But in some cases, we only need a subspace of \mathbb{C}^n as the state space of quantum information processing. This consideration motivates the following notions, which we call them quantum gate maximum and minimum subspace-fidelities, respectively. For a subspace $U \subseteq \mathbb{C}^n$,

$$F_U^{min}(\tilde{\mathbf{B}}) = \min_{u \in U} F_{\tilde{\mathbf{B}}, \tilde{I}}(u), \quad \text{and} \quad F_U^{max}(\tilde{\mathbf{B}}) = \max_{u \in U} F_{\tilde{\mathbf{B}}, \tilde{I}}(u).$$

Note that $F_U^{min}(\tilde{\mathbf{B}})$ and $F_U^{max}(\tilde{\mathbf{B}})$ quantify the worst-case and best-case behavior of the system by minimizing and maximizing over all possible initial states, respectively. Obviously, $0 \leq F_U^{min}(\tilde{\mathbf{B}}) \leq F_U^{max}(\tilde{\mathbf{B}}) \leq 1$ as $0 \leq F(\rho, \sigma) \leq 1$. We also have $F(\tilde{\mathbf{B}}) = \min_{U \subseteq \mathbb{C}^n} F_U^{min}(\tilde{\mathbf{B}}) = \min_{U \subseteq \mathbb{C}^n} F_U^{max}(\tilde{\mathbf{B}})$, where the second equation follows by examining one-dimensional subspaces.

One key notion in quantum error correction is that of noiseless subspaces, which have been intensively discussed in the setting where $\tilde{\mathbf{B}}$ is studied as a noise model [KL97, Lid14]. Intuitively, noiseless subspaces are shelters under quantum noise $\tilde{\mathbf{B}}$, as states in such a subspace are perfectly preserved under $\tilde{\mathbf{B}}$.

Definition 47. Let $\mathbf{B} = \{B_1, \dots, B_m\} \in \text{QC}(n, \mathbb{C})$. A noiseless subspace of \mathbf{B} is a non-zero subspace $U \subseteq \mathbb{C}^n$, such that for any $u \in U$, and any B_i , we have $B_i u = u$.

The following result formally shows that noiseless subspaces and isotropic subspaces are totally opposite from the viewpoint of the two quantum gate subspace-fidelities.

Proposition 48. *Let $\mathbf{B} = \{B_1, \dots, B_m\} \in \text{QC}(n, \mathbb{C})$.*

- $F_U^{\max}(\tilde{\mathbf{B}}) = 0$ if and only if U is an isotropic space;
- $F_U^{\max}(\tilde{\mathbf{B}}) = 1$ if and only if there is a noiseless subspace in U ;
- $F_U^{\min}(\tilde{\mathbf{B}}) = 1$ if and only if U is a noiseless subspace;
- $F_U^{\min}(\tilde{\mathbf{B}}) = 0$ if and only if there is an isotropic space in U .

Proof. These four claims are directly from the definitions of isotropic spaces, noiseless subspaces, quantum gate minimum subspace-fidelity and maximum subspace-fidelity. \square

Knill devised an efficient algorithm to find all noiseless subspaces for a given $\tilde{\mathbf{B}}$ [Kni06]. So we have a quite good understanding on $F_U^{\min}(\tilde{\mathbf{B}}) = 1$. On the other hand, isotropic subspaces fully characterize $F_U^{\max}(\tilde{\mathbf{B}}) = 0$. Therefore, isotropic spaces reveal the structure of the worst-case behavior of the channel.

Let us further point out another potential application of isotropic spaces in quantum control. A basic task of controlling quantum systems is to transfer all unknown quantum states into some targeting subspace [TV08, CT15]. So designing a control scheme as a quantum channel with a non-trivial isotropic space (the dimension greater than 1) can turn all quantum states residing in the isotropic space into the orthogonal complement of it.

Acknowledgement. We would like to thank Nengkun Yu, James B. Wilson, and Gábor Ivanyos for discussions related to this paper. We would also like to thank George Glauber and László Pyber for answering their questions on groups, including pointing out the reference [Ol'78] and clarifying the field characteristic issue in [BGH87]. Last but not least, we are grateful to the anonymous reviewers for their careful reading and very helpful feedback.

J. G. was partly supported by the National Key R&D Program of China (Grant No: 2018YFA0306701), the National Natural Science Foundation of China (Grant No: 61832015), and the Key Research Program of Frontier Sciences, CAS, Grant No. QYZDJ-SSW-SYS003. S. C. was partly supported by the National Key R&D Program of China (Grant No: 2018YFA0306701), and the National Natural Science Foundation of China (Grant No:61702489). Y. Q. was partly supported by the Australian Research Council DECRA DE150100720.

References

- [AB95] J. L. Alperin and R. B. Bell. *Groups and representations*. Number 162 in Graduate texts in mathematics. Springer, 1995.
- [Ada62] J. F. Adams. Vector fields on spheres. *Annals of Mathematics*, pages 603–632, 1962.
- [Alb39] A. A. Albert. *Structure of Algebras*. Number v. 24 in American Mathematical Society colloquium publications. American Mathematical Society, 1939.

- [ALP65a] J. F. Adams, Peter D. Lax, and Ralph S. Phillips. On matrices whose real linear combinations are nonsingular. *Proceedings of the American Mathematical Society*, 16(2):318–322, 1965.
- [Alp65b] J. L. Alperin. Large abelian subgroups of p -groups. *Transactions of the American Mathematical Society*, 117:10–20, 1965.
- [AM70] J. Gary Augustson and Jack Minker. An analysis of some graph theoretical cluster techniques. *Journal of the ACM*, 17(4):571–588, 1970.
- [Ant87] VA Antonov. Finite groups with a modular lattice of centralizers. *Algebra and Logic*, 26(6):403–422, 1987.
- [Bab85] László Babai. Trading group theory for randomness. In Robert Sedgewick, editor, *Proceedings of the 17th Annual ACM Symposium on Theory of Computing, May 6-8, 1985, Providence, Rhode Island, USA*, pages 421–429. ACM, 1985.
- [Bab16] László Babai. Graph isomorphism in quasipolynomial time [extended abstract]. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 684–697, 2016. arXiv:1512.03547, version 2.
- [Bae38] Reinhold Baer. Groups with abelian central quotient group. *Transactions of the American Mathematical Society*, 44(3):357–386, 1938.
- [BBS09] László Babai, Robert Beals, and Ákos Seress. Polynomial-time theory of matrix groups. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 55–64, 2009.
- [BCG⁺20] Xiaohui Bei, Shiteng Chen, Ji Guan, Youming Qiao, and Xiaoming Sun. From independent sets and vertex colorings to isotropic spaces and isotropic decompositions. arXiv:1904.03950; to appear in ITCS, 2020.
- [Bea95] Robert Beals. Towards polynomial time algorithms for matrix groups. In Larry Finkelstein and William M. Kantor, editors, *Groups and Computation, Proceedings of a DIMACS Workshop, New Brunswick, New Jersey, USA, June 7-10, 1995*, volume 28 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 31–54. DIMACS/AMS, 1995.
- [Ber84] Stuart J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Inf. Process. Lett.*, 18(3):147–150, 1984.
- [BFM13] Ada Boralevi, Daniele Faenzi, and Emilia Mezzetti. Linear spaces of matrices of constant rank and instanton bundles. *Advances in Mathematics*, 248:895–920, 2013.
- [BFS99] Jonathan F Buss, Gudmund S Frandsen, and Jeffrey O Shallit. The computational complexity of some problems of linear algebra. *Journal of Computer and System Sciences*, 58(3):572–596, 1999.
- [BGH87] Joe Buhler, Ranee Gupta, and Joe Harris. Isotropic subspaces for skewforms and maximal abelian subgroups of p -groups. *Journal of Algebra*, 108(1):269–279, 1987.

- [BH08] Andreas Björklund and Thore Husfeldt. Exact algorithms for exact satisfiability and number of perfect matchings. *Algorithmica*, 52(2):226–249, 2008.
- [BHK09] Andreas Björklund, Thore Husfeldt, and Mikko Koivisto. Set partitioning via inclusion-exclusion. *SIAM Journal on Computing*, 39(2):546–563, 2009.
- [BHK⁺16] Andreas Björklund, Thore Husfeldt, Petteri Kaski, Mikko Koivisto, Jesper Nederlof, and Pekka Parviainen. Fast zeta transforms for lattices with few irreducibles. *ACM Trans. Algorithms*, 12(1):4:1–4:19, 2016.
- [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.
- [Bol01] Béla Bollobás. *Random Graphs*. Number 73 in Cambridge Studies in Advanced Mathematics. Cambridge University Press, second edition, 2001.
- [BS84] László Babai and Endre Szemerédi. On the complexity of matrix group problems I. In *25th Annual Symposium on Foundations of Computer Science, West Palm Beach, Florida, USA, 24-26 October 1984*, pages 229–240, 1984.
- [Bur13] W. Burnside. On some properties of groups whose orders are powers of primes. *Proceedings of the London Mathematical Society*, 2(1):225–245, 1913.
- [BW12] Peter A. Brooksbank and James B. Wilson. Computing isometry groups of Hermitian maps. *Trans. Amer. Math. Soc.*, 364:1975–1996, 2012.
- [Bys04] Jesper Makhholm Byskov. Enumerating maximal independent sets with applications to graph colouring. *Operations Research Letters*, 32(6):547–556, 2004.
- [Cai94] Jin-yi Cai. Computing Jordan normal forms exactly for commuting matrices in polynomial time. *Int. J. Found. Comput. Sci.*, 5(3/4):293–302, 1994.
- [Can88] John F. Canny. Some algebraic and geometric computations in PSPACE. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*, pages 460–467, 1988.
- [Cat91] Fabrizio Catanese. Moduli and classification of irregular kaehler manifolds (and algebraic varieties) with albanese general type fibrations. *Inventiones mathematicae*, 104(1):263–289, 1991.
- [Chi85] Alexander L. Chistov. Fast parallel calculation of the rank of matrices over a field of arbitrary characteristic. In Lothar Budach, editor, *Fundamentals of Computation Theory, FCT '85, Cottbus, GDR, September 9-13, 1985*, volume 199 of *Lecture Notes in Computer Science*, pages 63–69. Springer, 1985.
- [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.

- [CKSU05] Henry Cohn, Robert D. Kleinberg, Balázs Szegedy, and Christopher Umans. Group-theoretic algorithms for matrix multiplication. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), 23-25 October 2005, Pittsburgh, PA, USA, Proceedings*, pages 379–388. IEEE Computer Society, 2005.
- [Coh75] P. M. Cohn. The word problem for free fields: A correction and an addendum. *J. Symbolic Logic*, 40(1):69–74, 03 1975.
- [CP07] Andrea Causin and Gian Pietro Pirola. A note on spaces of symmetric matrices. *Linear Algebra and Its Applications*, 2(426):533–539, 2007.
- [CR99] P. M. Cohn and C. Reutenauer. On the construction of the free field. *International Journal of Algebra and Computation*, 9(3-4):307–323, 1999.
- [CT15] Giuseppe Ilario Cirillo and Francesco Ticozzi. Decompositions of Hilbert spaces, stability analysis and convergence probabilities for discrete-time quantum dynamical semigroups. *Journal of Physics A: Mathematical and Theoretical*, 48(8):085302, 2015.
- [Dav70] E. B. Davies. Quantum stochastic processes ii. *Communications in Mathematical Physics*, 19(2):83–105, 1970.
- [Der01] H. Derksen. Polynomial bounds for rings of invariants. *Proceedings of the American Mathematical Society*, 129(4):955–964, 2001.
- [Die17] Reinhard Diestel. *Graph Theory*. Number 173 in Springer Graduate Texts in Mathematics. Springer, 5th edition, 2017.
- [Dim08] Alexandru Dimca. On the isotropic subspace theorems. *Bulletin mathématique de la Société des Sciences Mathématiques de Roumanie*, pages 307–324, 2008.
- [Dix71] John D. Dixon. Maximal abelian subgroups of the symmetric groups. *Canadian Journal of Mathematics*, 23(3):426–438, 1971.
- [DM17] H. Derksen and V. Makam. Polynomial degree bounds for matrix semi-invariants. *Advances in Mathematics*, 310:44–63, 2017.
- [dSP13] Clément de Seguins Pazzis. Large affine spaces of non-singular matrices. *Transactions of the American Mathematical Society*, 365(5):2569–2596, 2013.
- [Ebe91] Wayne Eberly. Decompositions of algebras over \mathbb{R} and \mathbb{C} . *Computational Complexity*, 1:211–234, 1991.
- [Epp03] David Eppstein. Small maximal independent sets and faster exact graph coloring. *J. Graph Algorithms Appl.*, 7(2):131–140, 2003.
- [FGT16] Stephen A. Fenner, Rohit Gurjar, and Thomas Thierauf. Bipartite perfect matching is in quasi-nc. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 754–763, 2016.

- [FP09] Franco Fagnola and Rely Pellicer. Irreducible and periodic positive maps. *Commun. Stoch. Anal*, 3(3):407–418, 2009.
- [FR04] M. Fortin and C. Reutenauer. Commutative/noncommutative rank of linear matrices and subspaces of matrices of low rank. *Séminaire Lotharingien de Combinatoire*, 52:B52f, 2004.
- [Gel17] I. Gelbukh. Isotropy index for the connected sum and the direct product of manifolds. *Publicationes Mathematicae Debrecen*, 90(3-4):287–310, 2017.
- [GFY18] Ji Guan, Yuan Feng, and Mingsheng Ying. Decomposition of quantum markov chains and its applications. *Journal of Computer and System Sciences*, 95:55–68, 2018.
- [GG07] Daniel Goldstein and Robert M. Guralnick. Alternating forms and self-adjoint operators. *Journal of Algebra*, 308(1):330–349, 2007.
- [GGdOW16] Ankit Garg, Leonid Gurvits, Rafael Mendes de 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.
- [GGdOW17] Ankit Garg, Leonid Gurvits, Rafael Mendes de Oliveira, and Avi Wigderson. Algorithmic and optimization aspects of brascamp-lieb inequalities, via operator scaling. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 397–409, 2017.
- [GL84] R. Gow and T. J. Laffey. Pairs of alternating forms and products of two skew-symmetric matrices. *Linear algebra and its applications*, 63:119–132, 1984.
- [Gur04] Leonid Gurvits. Classical complexity and quantum entanglement. *J. Comput. Syst. Sci.*, 69(3):448–484, 2004.
- [Hal35] P. Hall. On representatives of subsets. *Journal of the London Mathematical Society*, 1(1):26–30, 1935.
- [HW15] Pavel Hrubeš and Avi Wigderson. Non-commutative arithmetic circuits with division. *Theory of Computing*, 11:357–393, 2015.
- [IKQS15] Gábor Ivanyos, Marek Karpinski, Youming Qiao, and Miklos Santha. Generalized Wong sequences and their applications to Edmonds’ problems. *J. Comput. Syst. Sci.*, 81(7):1373–1386, 2015.
- [IL99] Bo Ilic and J. M. Landsberg. On symmetric degeneracy loci, spaces of symmetric matrices of constant rank and dual varieties. *Mathematische Annalen*, 314(1):159–174, 1999.
- [IQ19] 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. *SIAM Journal on Computing*, 48(3):926–963, 2019.

- [IQS17] Gábor Ivanyos, Youming Qiao, and K. V. Subrahmanyam. Non-commutative Edmonds’ problem and matrix semi-invariants. *Computational Complexity*, 26(3):717–763, 2017.
- [IQS18] Gábor Ivanyos, Youming Qiao, and K. V. Subrahmanyam. Constructive non-commutative rank computation is in deterministic polynomial time. *Computational Complexity*, 27(4):561–593, 2018.
- [IR99] Gábor Ivanyos and Lajos Rónyai. Computations in associative and Lie algebras. In *Some tapas of computer algebra*, pages 91–120. Springer, 1999.
- [JPY88] David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. On generating all maximal independent sets. *Inf. Process. Lett.*, 27(3):119–123, 1988.
- [JT95] Tommy R. Jensen and Bjarne Toft. *Graph coloring problems*, volume 39 of *Wiley-Interscience series in discrete mathematics and optimization*. John Wiley & Sons, 1995.
- [KI04] Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity*, 13(1-2):1–46, 2004.
- [KL97] Emanuel Knill and Raymond Laflamme. Theory of quantum error-correcting codes. *Physical Review A*, 55(2):900, 1997.
- [Kni06] Emanuel Knill. Protected realizations of quantum information. *Physical Review A*, 74(4):042301, 2006.
- [Koi96] Pascal Koiran. Hilbert’s nullstellensatz is in the polynomial hierarchy. *J. Complexity*, 12(4):273–286, 1996.
- [KUW86] Richard M. Karp, Eli Upfal, and Avi Wigderson. Constructing a perfect matching is in random NC. *Combinatorica*, 6(1):35–48, 1986.
- [Lan02] Serge Lang. *Algebra*. Number 211 in Graduate Texts in Mathematics. Springer-Verlag, New York, third enlarged edition, 2002.
- [Law76] Eugene L. Lawler. A note on the complexity of the chromatic number problem. *Inf. Proc. Lett.*, 5:66–67, 1976.
- [Lid14] Daniel A Lidar. Review of decoherence free subspaces, noiseless subsystems, and dynamical decoupling. *Adv. Chem. Phys.*, 154:295–354, 2014.
- [LLK80] Eugene L. Lawler, Jan Karel Lenstra, and A. H. G. Rinnooy Kan. Generating all maximal independent sets: Np-hardness and polynomial-time algorithms. *SIAM J. Comput.*, 9(3):558–565, 1980.
- [Lov79] László Lovász. On determinants, matchings, and random algorithms. In *FCT*, pages 565–574, 1979.

- [Lov86] László Lovász. *An Algorithmic Theory of Numbers, Graphs, and Convexity*, volume 50 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. SIAM, 1986.
- [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.
- [LP17] David A. Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.
- [LQ17] Yinan Li and Youming Qiao. Linear algebraic analogues of the graph isomorphism problem and the Erdős-Rényi model. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 463–474, 2017. arXiv:1708.04501, v2.
- [LSW00] Nathan Linial, Alex Samorodnitsky, and Avi Wigderson. A deterministic strongly polynomial algorithm for matrix scaling and approximate permanents. *Combinatorica*, 20(4):545–568, 2000.
- [LW12] Mark L. Lewis and James B. Wilson. Isomorphism in expanding families of indistinguishable groups. *Groups - Complexity - Cryptology*, 4(1):73–110, 2012.
- [LY93] Kee Yuen Lam and Paul Yiu. Linear spaces of real matrices of constant rank. *Linear algebra and its applications*, 195:69–79, 1993.
- [Mes90] Roy Meshulam. On k -spaces of real matrices. *Linear and multilinear algebra*, 26(1-2):39–41, 1990.
- [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.
- [MM65] John W. Moon and Leo Moser. On cliques in graphs. *Israel Journal of Mathematics*, 3(1):23–28, 1965.
- [Mul12] Ketan Mulmuley. Geometric complexity theory V: equivalence between blackbox derandomization of polynomial identity testing and derandomization of noether’s normalization lemma. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 629–638, 2012.
- [Mul17] Ketan Mulmuley. Geometric complexity theory V: Efficient algorithms for Noether normalization. *Journal of the American Mathematical Society*, 30(1):225–309, 2017.
- [MVV87] Ketan Mulmuley, Umesh V. Vazirani, and Vijay V. Vazirani. Matching is as easy as matrix inversion. *Combinatorica*, 7(1):105–113, 1987.
- [NC02] Michael A Nielsen and Isaac Chuang. *Quantum computation and quantum information*, 2002.

- [Ol'78] A. Yu Ol'shanskii. The number of generators and orders of abelian subgroups of finite p -groups. *Mathematical notes of the Academy of Sciences of the USSR*, 23(3):183–185, 1978.
- [Qia19] Youming Qiao. Matrix spaces as a linear algebraic analogue of graphs. Under preparation, 2019.
- [Ren92] James Renegar. On the computational complexity and geometry of the first-order theory of the reals, part I: introduction. preliminaries. the geometry of semi-algebraic sets. the decision problem for the existential theory of the reals. *J. Symb. Comput.*, 13(3):255–300, 1992.
- [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.
- [Rón94] Lajos Rónyai. A deterministic method for computing splitting elements in simple algebras over \mathbb{Q} . *J. Algorithms*, 16(1):24–32, 1994.
- [Ros11] Tobias Rossmann. *Algorithms for Nilpotent Linear Groups*. PhD thesis, National University of Ireland, Galway, 2011.
- [Sch76] Rudolf Scharlau. Pairs of alternating forms. *Mathematische Zeitschrift*, 147(1):13–19, 1976.
- [Sch80] Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, 1980.
- [She11] John Sheekey. *On rank problems for subspaces of matrices over finite fields*. PhD thesis, University College Dublin, 2011.
- [Sin64] Richard Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *The annals of mathematical statistics*, 35(2):876–879, 1964.
- [ST17] Ola Svensson and Jakub Tarnawski. The matching problem in general graphs is in quasi-nc. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 696–707. IEEE Computer Society, 2017.
- [Syl86] John Sylvester. On the dimension of spaces of linear transformations satisfying rank conditions. *Linear Algebra and its Applications*, 78:1–10, 1986.
- [TIAS77] Shuji Tsukiyama, Mikio Ide, Hiromu Ariyoshi, and Isao Shirakawa. A new algorithm for generating all the maximal independent sets. *SIAM Journal on Computing*, 6(3):505–517, 1977.
- [Tur41] Paul Turán. Egy gráfelméleti szélsőértékfeladatról (on an extremal problem in graph theory). *Mat. Fiz. Lapok*, 48:436–452, 1941.

- [Tut47] William T. Tutte. The factorization of linear graphs. *Journal of the London Mathematical Society*, 1(2):107–111, 1947.
- [TV08] Francesco Ticozzi and Lorenza Viola. Quantum markovian subsystems: invariance, attractivity, and control. *IEEE Transactions on Automatic Control*, 53(9):2048–2063, 2008.
- [Val79] Leslie G. Valiant. Completeness classes in algebra. In *Proceedings of the 11th Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1979, Atlanta, Georgia, USA*, pages 249–261, 1979.
- [Vat11] Vincent Vatter. Maximal independent sets and separating covers. *The American Mathematical Monthly*, 118(5):418–423, 2011.
- [Vdo04] EP Vdovin. The number of subgroups with trivial unipotent radicals in finite groups of lie type. *Journal of Group Theory*, 7(1):99–112, 2004.
- [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.
- [Wes87] R. Westwick. Spaces of matrices of fixed rank. *Linear and Multilinear Algebra*, 20(2):171–174, 1987.
- [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.
- [Wil13] James B. Wilson. Optimal algorithms of Gram–Schmidt type. *Linear Algebra and its Applications*, 438(12):4573–4583, 2013.
- [Wol12] Michael M. Wolf. Quantum channels & operations: Guided tour, May 2012. Lecture notes available at <https://www-m5.ma.tum.de/foswiki/pub/M5/Allgemeines/MichaelWolf/QChannelLecture.pdf>.
- [Woo89] Jay A Wood. Spinor groups and algebraic coding theory. *Journal of Combinatorial Theory, Series A*, 51(2):277–313, 1989.
- [Woo11] David Wood. On the number of maximal independent sets in a graph. *Discrete Mathematics and Theoretical Computer Science*, 13(3):17, 2011.
- [Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In Edward W. Ng, editor, *Symbolic and Algebraic Computation, EUROSAM '79, An International Symposium on Symbolic and Algebraic Computation, Marseille, France, June 1979, Proceedings*, volume 72 of *Lecture Notes in Computer Science*, pages 216–226. Springer, 1979.