

Connections between graphs and matrix spaces

Yinan Li* Youming Qiao[†] Avi Wigderson[‡] Yuval Wigderson[§] Chuanqi Zhang[¶]

January 5, 2023

Abstract

Given a bipartite graph G , the graphical matrix space \mathcal{S}_G consists of matrices whose non-zero entries can only be at those positions corresponding to edges in G . Tutte (*J. London Math. Soc.*, 1947), Edmonds (*J. Res. Nat. Bur. Standards Sect. B*, 1967) and Lovász (*FCT*, 1979) observed connections between perfect matchings in G and full-rank matrices in \mathcal{S}_G . Dieudonné (*Arch. Math.*, 1948) proved a tight upper bound on the dimensions of those matrix spaces containing only singular matrices. The starting point of this paper is a simultaneous generalization of these two classical results: we show that the largest dimension over subspaces of \mathcal{S}_G containing only singular matrices is equal to the maximum size over subgraphs of G without perfect matchings, based on Meshulam’s proof of Dieudonné’s result (*Quart. J. Math.*, 1985).

Starting from this result, we go on to establish more connections between properties of graphs and matrix spaces. For example, we establish connections between acyclicity and nilpotency, between strong connectivity and irreducibility, and between isomorphism and conjugacy/congruence. For each connection, we study three types of correspondences, namely the basic correspondence, the inherited correspondence (for subgraphs and subspaces), and the induced correspondence (for induced subgraphs and restrictions). Some correspondences lead to intriguing generalizations of classical results, such as Dieudonné’s result mentioned above, and a celebrated theorem of Gerstenhaber regarding the largest dimension of nil matrix spaces (*Amer. J. Math.*, 1958).

Finally, we show some implications of our results to quantum information and present open problems in computational complexity motivated by these results.

*School of Mathematics and Statistics, Wuhan University, China (Yinan.Li@whu.edu.cn). Research supported in part by MEXT Quantum Leap Flagship Program (MEXT Q-LEAP) Grant Number JPMXS0120319794.

[†]Centre for Quantum Software and Information, University of Technology Sydney, Australia (Youming.Qiao@uts.edu.au). Research supported in part by Australian Research Council DP200100950.

[‡]School of Mathematics, Institute for Advanced Study, Princeton, New Jersey 08540 (avi@ias.edu). Research supported in part by NSF grant CCF-1900460.

[§]Department of Mathematics, Tel Aviv University, Israel (yuvalwig@tauex.tau.ac.il). Research supported in part by NSF GRFP Grant DGE-1656518, ERC Consolidator Grant 863438, ERC Starting Grant 101044123 and NSF-BSF Grant 20196.

[¶]Centre for Quantum Software and Information, University of Technology Sydney, Australia (Chuanqi.Zhang@student.uts.edu.au). Research supported by Australian Research Council DP200100950 and the Sydney Quantum Academy, Sydney, NSW, Australia.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 3 |
| 1.1 | Overview | 3 |
| 1.2 | Main results | 6 |
| 1.2.1 | Matchings in bipartite graphs and ranks of matrices | 6 |
| 1.2.2 | Cycles in directed graphs and nilpotent matrices | 7 |
| 1.2.3 | Strong connectivity and irreducibility | 8 |
| 1.2.4 | Isomorphism, conjugacy, and congruence | 9 |
| 1.2.5 | Vertex transitivity and conjugacy/congruence irreducibility | 10 |
| 1.3 | Perspective: symmetries of matrix spaces | 10 |
| 1.4 | Perspective: computational complexity | 12 |
| 1.5 | Connections with quantum information theory | 14 |
| 1.6 | Some known connections between graphs and matrix spaces | 15 |
| 1.7 | Discussion, research directions, and open problems | 16 |
| 2 | Preliminaries | 17 |
| 3 | Matchings in bipartite graphs and ranks of matrices | 17 |
| 3.1 | The inherited correspondence | 17 |
| 3.2 | The induced correspondence | 18 |
| 4 | Cycles in directed graphs and nilpotent matrices | 19 |
| 4.1 | The basic correspondences | 19 |
| 4.1.1 | Maximum walk lengths and nil/nilpotent indices | 19 |
| 4.1.2 | Cycle covers and the number of zero eigenvalues | 21 |
| 4.2 | The inherited correspondence | 22 |
| 4.3 | The induced correspondence | 24 |
| 5 | Strong connectivity and irreducibility | 25 |
| 5.1 | The basic correspondences | 25 |
| 5.2 | The inherited correspondence | 26 |
| 5.3 | The induced correspondence | 27 |
| 6 | Isomorphism, congruence, and conjugacy | 28 |
| 6.1 | Isomorphism and congruence | 28 |
| 6.2 | Isomorphism and conjugacy | 29 |
| 7 | Vertex transitivity and conjugacy/congruence irreducibility | 32 |
| 8 | Connections to quantum information theory | 35 |
| 8.1 | Quantum channels and transition matrices | 35 |
| 8.2 | Quantum expanders and dimension expanders | 37 |
| 8.3 | Quantum channels and zero-error communication | 38 |
| A | On finding cycles in directed graphs | 44 |

1 Introduction

1.1 Overview

Let $M(n, \mathbb{F})$ denote the vector space of $n \times n$ square matrices over a field \mathbb{F} . A vector subspace \mathcal{S} of $M(n, \mathbb{F})$ is called a *matrix space*. Matrix spaces are basic and fundamental mathematical objects, and they arise naturally in many different areas of mathematics, physics, and computer science. In algebraic geometry, they arise in close connection with certain sheaves on projective space [EH88]. In topology, they arise naturally in connection to linearly independent vector fields on spheres, which led to the development of the Adams operations on topological K -theory [Ada62, ALP65]. In invariant theory, they were used by Dieudonné [Die48] to classify the symmetries of the determinant, recovering a result of Frobenius [Fro97]. Gerstenhaber [Ger58] used matrix spaces to make progress on Albert’s problem in the theory of non-associative algebras. In group theory, Baer observes that (alternating) matrix spaces are closely connected with p -groups of class 2 [Bae38]. Recently, they arise in the study of completely positive maps and quantum expanders in quantum information theory [BBLM20]. In computational complexity theory, they underlie the *polynomial identity testing* problem, a central challenge in derandomization and algebraic complexity [KI04]. Finally, in discrete mathematics and theoretical computer science, matrix spaces can be used to study matchings in graphs, and this connection has important algorithmic consequences [Tut47, Edm67, Lov79, Lov89]. We return to this example shortly.

One important way of thinking of matrix spaces is as *symbolic matrices*. Namely, we may choose a basis for \mathcal{S} and represent a generic element of \mathcal{S} as a generic linear combination of the basis elements. By doing so, we construct a matrix whose entries are homogeneous linear forms in some variables, and we can recover \mathcal{S} by substituting all possible elements of \mathbb{F} into these variables. We discuss symbolic matrices in more detail in Section 1.4.

In this paper, we focus on matrix spaces of *restricted support*. Namely, we fix a set of positions $E \subseteq [n]^2$, and study properties of matrix spaces $\mathcal{S} \leq M(n, \mathbb{F})$ satisfying that every matrix $M \in \mathcal{S}$ is supported on E , namely that the (i, j) th entry of M is equal to zero for all $(i, j) \notin E$. Crucially for our purposes, we can encode the restricted support information as a *graph*. Indeed, the support set $E \subseteq [n] \times [n]$ is a set of ordered pairs, and this can naturally be viewed as either the arc set of a directed graph on the vertex set $[n]$ or as the edge set of a bipartite graph with vertex set $[n] \times [n]$. Formally, we can make the following definition.

Definition 1.1 (Graphical matrix spaces). For $(i, j) \in [n]^2$, let $E_{i,j}$ be the elementary matrix in $M(n, \mathbb{F})$ where the (i, j) th entry is 1, and the remaining entries are 0. Suppose $G = (L \cup R, E)$ is a bipartite graph, where $L = R = [n]$, or $G = ([n], E)$ is a directed graph. The *graphical matrix space* \mathcal{S}_G (over \mathbb{F}) corresponding to G the subspace of $M(n, \mathbb{F})$ spanned by $\{E_{i,j} \mid (i, j) \in E\}$.

Thus, we see that a matrix space supported on the edges of G is the same as a subspace of \mathcal{S}_G . Additionally, we may view the symbolic matrix defining some $\mathcal{S} \leq \mathcal{S}_G$ as simply a collection of homogeneous linear forms, each of which is associated with an edge of G .

At first glance, it is not clear why encoding the restricted support as a graph is helpful or meaningful. However, it turns out that many natural linear-algebraic properties of matrix spaces correspond directly to graph-theoretic properties of (bipartite or directed) graphs, and this connection is captured by the association of the matrix space \mathcal{S}_G to the graph G . Moreover, the study of matrix spaces with restricted support—that is, subspaces of \mathcal{S}_G —closely mirrors the study of subgraphs of G . Many specific connections of this type were studied in the past, as we discuss below. In this paper, we initiate a systematic study of the connections between graphs and their associated graphical matrix spaces. In many instances, these explorations yield surprising generalizations and extensions of known results in the “full-support” setting.

To motivate our results, we begin with a connection mentioned above, namely the connection between perfect matchings in bipartite graphs and singularity of matrix spaces, beginning with foundational works of Tutte [Tut47], Edmonds [Edm67], and Lovász [Lov79]. Let G be a bipartite graph with both parts of size n . The key observation is that a linear-algebraic property of the matrix space \mathcal{S}_G encodes a graph-theoretic property of G . Namely, G has a perfect matching if and only if \mathcal{S}_G contains a non-singular matrix. Indeed, if G has a perfect matching, then the matrix in \mathcal{S}_G which has a 1 on every edge of this perfect matching and zeroes elsewhere, is non-singular. Conversely, if \mathcal{S}_G contains a non-singular matrix, then the expansion of its determinant has at least one non-zero summand, which corresponds to a perfect matching in G . Equivalently, if we encode the matrix space \mathcal{S}_G as a symbolic matrix, then G has a perfect matching if and only if the determinant of this symbolic matrix is not the zero polynomial.

The first important consequence of this connection is algorithmic. Indeed, Lovász [Lov79] used this connection to find an RNC algorithm¹ for the problem of determining whether a bipartite graph has a perfect matching. Additionally, inspired by this connection, Edmonds [Edm67] asked whether one could devise an efficient, *deterministic* algorithm to determine if such an arbitrary matrix space contains a non-singular matrix. This is one of the most important problems in computational complexity, and we return to it in Section 1.4.

In this paper, we extend the basic connection between perfect matchings in graphs and non-singular matrices in matrix spaces in several ways. First, we develop a number of other correspondences, showing that a graph (or directed graph) has some graph-theoretic property if and only if its associated graphical matrix space has some linear-algebraic property. In most instances this connection is fairly straightforward to prove (like in the case above), but in some others it is quite involved and requires a number of different ideas. Secondly, we prove a number of *dimension theorems*, which demonstrate that in some instances, these correspondences between graphs and their graphical matrix spaces extend to their subgraphs and subspaces, respectively, and are surprisingly deep. Third, we prove a number of *order theorems*, which demonstrate that in some instances, these correspondences between graphs and their graphical matrix spaces extend to their induced subgraphs and induced subspaces, respectively.

Continuing the discussion above, here is an example of one of our dimension theorems, which is a special case of Theorem 1.5.

Theorem 1.2. *Let G be a bipartite graph with two parts of size n , and let $\mathcal{S}_G \leq M(n, \mathbb{F})$ be the associated matrix space. The maximum number of edges of a subgraph of G with no perfect matching equals the largest dimension of a subspace of \mathcal{S}_G containing only singular matrices.*

Note that if H is an m -edge subgraph of G with no perfect matching, then $\mathcal{S}_H \leq \mathcal{S}_G$ is an m -dimensional matrix space supported on the edges of G , and every matrix in \mathcal{S}_H is singular. Theorem 1.2 says that such examples are extremal in the sense of dimension: if we wish to construct a singular matrix space \mathcal{S} supported on the edges of G , then the biggest \mathcal{S} we can take is “axis-aligned”, i.e., of the form \mathcal{S}_H for some $H \subseteq G$ with no perfect matching.

Theorem 1.2 generalizes a famous theorem of Dieudonné² [Die48], who proved that if every matrix in $\mathcal{S} \leq M(n, \mathbb{F})$ is singular, then $\dim \mathcal{S} \leq n(n-1)$. This is the special case of Theorem 1.2 in which $G = K_{n,n}$, as it is easy to check that the largest subgraph of $K_{n,n}$ with no perfect matching has $n(n-1)$ edges. In other words, Theorem 1.2 is simply the “restricted support” version of Dieudonné’s theorem: it determines the largest dimension of a singular matrix space with (any)

¹Informally, an efficient, probabilistic parallel algorithm.

²Dieudonné was interested in classifying the symmetries of the determinant, and used this result to find a new proof of Frobenius’s [Fro97] characterization of these symmetries.

restricted support, just as Dieudonné determined the largest dimension of a singular matrix space (with no restriction on its support). Our proof of Theorem 1.2 is based on Meshulam’s [Mes85] proof of Dieudonné’s theorem, and indeed Theorem 1.2 is almost implicit in Meshulam’s work.

We also present an example of one of our order theorems, which is a special case of Theorem 1.7. For this, we need the following notion of induced matrices. Let $B \in M(n, \mathbb{F})$ and $U, V \leq \mathbb{F}^n$ be dimension- s and dimension- t subspaces of \mathbb{F}^n , respectively. Viewing B as a bilinear form $\mathbb{F}^n \times \mathbb{F}^n \rightarrow \mathbb{F}$, the induced matrix $B[U, V] \in M(s \times t, \mathbb{F})$ with respect to $U, V \leq \mathbb{F}^n$ is obtained by restricting the first argument to U and the second to V . The *induced subspace* $\mathcal{S}[U, V]$ of $\mathcal{S} \leq M(n, \mathbb{F})$ with respect to $U, V \leq \mathbb{F}^n$ is the matrix space consisting of $B[U, V]$ for all $B \in \mathcal{S}$. The *order* of $\mathcal{S}[U, V]$ is $\dim(U) + \dim(V) = s + t$.

Theorem 1.3. *Let G be a bipartite graph with two parts of size n , and let $\mathcal{S}_G \leq M(n, \mathbb{F})$ be the associated matrix space. The maximum number of vertices of an induced subgraph of G with no perfect matching equals the largest order of an induced subspace of \mathcal{S}_G containing only singular matrices.*

Most of our results fit into a similar framework to what was described above. First, we have what we call a *basic correspondence* between a graph-theoretic property P and a linear-algebraic property Q , which simply means that a graph G has property P if and only if its graphical matrix space \mathcal{S}_G has property Q . The dimension theorem boosts this to what we call an *inherited correspondence*: this says that the maximum number of edges in a spanning subgraph $H \subseteq G$ with property P is equal to the largest dimension of a matrix space $\mathcal{S} \leq \mathcal{S}_G$ with property Q . Similarly, the order theorem boosts the basic correspondence to what we call an *induced correspondence*: in this case, the number of vertices in the maximum induced subgraph H of G with property P equals the order of the maximum induced subspace of \mathcal{S}_G with property Q .

One useful way of thinking about inherited correspondences like Theorem 1.2 is in terms of the symmetries at play. The natural group action on the set of bipartite graphs with both vertex sets of size n is that of the group $S_n \times S_n$, which acts by simultaneously permuting the vertices in each of the two parts. This corresponds to relabeling the vertices in each part, and it of course preserves basic graph-theoretic properties such as the existence of a perfect matching. Similarly, in this context, a natural action on the space $M(n, \mathbb{F})$ of $n \times n$ matrices is the left-right action of $GL(n, \mathbb{F}) \times GL(n, \mathbb{F})$, where we multiply on the left and the right by two invertible matrices. If we view matrices in $M(n, \mathbb{F})$ as linear maps $\mathbb{F}^n \rightarrow \mathbb{F}^n$, then this action corresponds to simultaneously changing bases in the domain and codomain. This action preserves many important linear-algebraic properties of $n \times n$ matrices, such as singularity. Of course, $S_n \times S_n$ is a subgroup of $GL(n, \mathbb{F}) \times GL(n, \mathbb{F})$, which simply represents the fact that permuting the rows and columns of a matrix is one way of changing bases on the domain and codomain.

However, $GL(n, \mathbb{F}) \times GL(n, \mathbb{F})$ is a much larger and richer group than $S_n \times S_n$, and its action on $M(n, \mathbb{F})$ is much more complicated than simply permuting the rows and columns. In the same way, the lattice of subspaces of \mathcal{S}_G is much richer than the lattice of subgraphs of G . Nonetheless, Theorem 1.2, as well as our other dimension theorems, say that in some instances, this additional structure adds no extra flexibility: the largest dimension of a subspace of \mathcal{S}_G with some property Q equals the largest dimension of “axis-aligned” subspaces \mathcal{S}_G with property Q , even though the set of such subspaces is much poorer.

In the rest of this introduction, we discuss our results in further detail, and explain some more about their connections to other topics. In Section 1.2, we state our main results; this is arranged into a collection of subsections, with each one discussing the correspondences between some graph-theoretic and linear-algebraic properties. In Section 1.3, we discuss the importance of underlying group actions to our results. In Sections 1.4 and 1.5, we discuss the connections to

computational complexity and quantum information theory, respectively. In Section 1.6, we broaden our scope, and mention other results which are of the form we discuss, namely correspondences between graphs and matrix spaces. Finally, in Section 1.7, we discuss future research directions and open problems.

1.2 Main results

We now describe our results in more detail. We begin with three topics in which we have dimension theorems: matchings in bipartite graphs, cycles in directed graphs, and strong connectivity of directed graphs. We then discuss several instances where we can prove a basic correspondence but not an inherited correspondence (and in some instances, we can even prove that inherited correspondence cannot hold). We first introduce some basic notation that we will use (and have used).

For $n \in \mathbb{N}$, $[n] := \{1, 2, \dots, n\}$. Let \mathbb{F}^n be the linear space of length- n row vectors over \mathbb{F} . Let $M(n, \mathbb{F})$ (resp. $M(m \times n, \mathbb{F})$) be the linear space of $n \times n$ (resp. $m \times n$) matrices over a field \mathbb{F} . For a matrix $B \in M(n, \mathbb{F})$ and a vector $v \in \mathbb{F}^n$, B acts on v from the right: vB is another row vector in \mathbb{F}^n . For a subspace $U \leq \mathbb{F}^n$, let $B(U) = \{uB : u \in U\} \leq \mathbb{F}^n$. Let $\mathcal{S} = \langle B_1, \dots, B_m \rangle$ be the linear span of matrices $B_1, \dots, B_m \in M(n, \mathbb{F})$. We call \mathcal{S} a *matrix space*, denoted as $\mathcal{S} \leq M(n, \mathbb{F})$. Let $GL(n, \mathbb{F})$ be the group of $n \times n$ invertible matrices over \mathbb{F} . For a matrix $T \in GL(n, \mathbb{F})$, let T^{-1} be its inverse. For a matrix $B \in M(n, \mathbb{F})$, let B^t be its transpose. If $\mathbb{F} = \mathbb{C}$, let B^* be its conjugate transpose.

A directed graph is $G = (V, E)$, where V is the vertex set and $E \subseteq V \times V$ is the *arc* set. We shall mostly work with directed graphs with vertex sets being $V = [n]$.

A bipartite graph is $G = (L \cup R, E)$, where L and R are the left and right vertex sets, and $E \subseteq L \times R$ is the *edge* set.

Except where otherwise stated, \mathbb{F} is an arbitrary field; except where explicitly specified, our results hold for all fields.

1.2.1 Matchings in bipartite graphs and ranks of matrices

Given a bipartite graph $G = ([m] \cup [n], E)$ with vertex parts $[m]$ and $[n]$, we let $\mathcal{S}_G = \langle E_{i,j} \mid (i, j) \in E \rangle \leq M(m \times n, \mathbb{F})$ be the graphical matrix space corresponding to G . A *matching* is a subset of edges where any two edges do not share common vertices. The *matching number* is the size of a maximum matching, i.e., a matching containing the largest possible number of edges. We begin by stating formally the basic correspondence which was discussed above, between perfect matchings and non-singular matrices.

Theorem 1.4 ([Edm67, Theorem 1]). *Let $G = ([m] \cup [n], E)$ be a bipartite graph with $m \leq n$ and $\mathcal{S}_G \leq M(m \times n, \mathbb{F})$ be the graphical matrix space associated with G . Then for each $r \in [m]$, the matching number of G is at most r if and only if the rank of every matrix in \mathcal{S}_G is at most r .*

Thus, in our terminology, Theorem 1.4 establishes a basic correspondence between the graph property P_r of the matching number being at most r and the linear-algebraic property Q_r of all matrices having rank at most r . Our next result, the more general form of Theorem 1.2, is the corresponding inherited correspondence between P_r and Q_r .

Theorem 1.5. *Let $G = ([m] \cup [n], E)$ be a bipartite graph with $m \leq n$ and $\mathcal{S}_G \leq M(m \times n, \mathbb{F})$ be the graphical matrix space associated with G . Then for each $r \in [m]$, the maximum size over subgraphs of G whose matching number is at most r equals the largest dimension over subspaces of \mathcal{S}_G in which every matrix is of rank at most r .*

This theorem is proved in Section 3.1. The proof is based on Meshulam’s proof of the Dieudonné–Flanders–Meshulam theorem [Mes85], which gives a bound on the dimension of an $m \times n$ matrix space containing matrices of rank at most r . In fact, the Dieudonné–Flanders–Meshulam theorem corresponds to taking the complete bipartite graph $G = K_{m,n}$ in Theorem 1.5 .

In addition to the basic and inherited correspondences above, we prove one further correspondence between matchings and ranks in the context of *induced correspondences*. This is a more general version of Theorem 1.3. We first formally define what we mean an induced matrix space here.

Definition 1.6. Let $\mathcal{S} \leq M(m \times n, \mathbb{F})$ be a matrix space over the field \mathbb{F} . For a subspace $L \leq \mathbb{F}^m$ (resp. $R \leq \mathbb{F}^n$) of dimension s (resp. t), let T_L (resp. T_R) be an $s \times m$ (resp. $t \times n$) matrix whose rows span L (resp. R). The *induced subspace* of \mathcal{S} on L and R is defined as $\mathcal{S}[L, R] := \{T_L B T_R^t \mid B \in \mathcal{S}\} \leq M(s \times t, \mathbb{F})$.³ The *order* of $\mathcal{S}[L, R]$ is $s + t$.

We discuss in more detail why this is a natural definition in Section 1.3 and for now simply remark that it mimics the notion of an induced subgraph. The induced correspondence in this context is as follows, proved in Section 3.2.

Theorem 1.7. *Let $G = ([m] \cup [n], E)$ be a bipartite graph with $m \leq n$ and $\mathcal{S}_G \leq M(m \times n, \mathbb{F})$ be the graphical matrix space associated with G . Then for each $r \in [m]$, the maximum order over induced subgraphs of G whose matching number is at most r equals the maximum order of an induced subspace of \mathcal{S}_G in which every matrix is of rank at most r .*

1.2.2 Cycles in directed graphs and nilpotent matrices

Given a directed graph $G = ([n], E)$ with vertex set $[n]$, we let the corresponding graphical matrix space be $\mathcal{S}_G = \langle E_{i,j} \mid (i, j) \in E \rangle \leq M(n, \mathbb{F})$. Note that the natural group action on the set of directed graphs is that of S_n , which acts by permuting the vertices. In the linear-algebraic setting, we will be focusing on the conjugation action of $GL(n, \mathbb{F})$ on $M(n, \mathbb{F})$. If we identify $M(n, \mathbb{F})$ with the set of endomorphisms of \mathbb{F}^n , then this action simply corresponds to changing the basis of \mathbb{F}^n . We have that S_n is a subgroup of $GL(n, \mathbb{F})$, corresponding to the fact that we may simultaneously permute the rows and columns of a matrix by a change of basis.

Recall that a directed graph is called *acyclic* if it contains no directed cycle. We also say that a matrix space $S \leq M(n, \mathbb{F})$ is *nil*⁴ if every matrix in S is nilpotent. We prove a number of correspondences between the graph-theoretic property of being acyclic and the linear-algebraic property of being nil, beginning with the following basic correspondence.

Theorem 1.8 (Simplified version of Theorem 4.2). *Let $G = ([n], E)$ be a directed graph and $\mathcal{S}_G \leq M(n, \mathbb{F})$ be the graphical matrix space associated with G . Then G is acyclic if and only if \mathcal{S}_G is nil.*

We are also able to prove more refined basic correspondences which generalize Theorem 1.8. First, we prove a basic correspondence between the maximum path length in G and the maximum nilpotent index of a matrix in \mathcal{S}_G ; see Theorem 4.2. Second, we prove a basic correspondence

³While different bases lead to different subspaces of $M(s \times t, \mathbb{F})$, these subspaces are unique up to equivalence. More precisely, let T'_L (resp. T'_R) be another $s \times m$ (resp. $t \times n$) matrix whose rows span L (resp. R) and let $\mathcal{S}[L, R]' := \{T'_L B (T'_R)^t \mid B \in \mathcal{S}\} \leq M(s \times t, \mathbb{F})$, we can find invertible matrices $A_L \in GL(s, \mathbb{F})$ and $A_R \in GL(t, \mathbb{F})$ such that $A_L T_L = T'_L$ and $A_R T_R = T'_R$. Thus $A_L \mathcal{S}[L, R] A_R^t = \mathcal{S}[L, R]'$.

⁴We adopt the terminology *nil* here following the practice in algebra, where the distinction of nil and nilpotent algebras naturally leads to the definitions of nil and nilpotent matrix spaces; see Section 4.

between the maximum cycle cover size in G and the maximum number of zero eigenvalues of a matrix in \mathcal{S}_G ; see Theorem 4.4.

Moreover, we prove a dimension version of Theorem 1.8, i.e., to boost the basic correspondence to an inherited correspondence, as stated in the following theorem.

Theorem 1.9. *Let $G = ([n], E)$ be a directed graph and $\mathcal{S}_G \leq M(n, \mathbb{F})$ be the graphical matrix space associated with G . The maximum number of arcs in an acyclic subgraph of G equals the largest dimension of a nil subspace of \mathcal{S}_G .*

See Section 4.2 for its proof. Theorem 1.9 is a generalization of a well-known theorem of Gerstenhaber about the largest dimension of a nil matrix spaces [Ger58] (cf. Remark 4.8), and its proof is adapted from de Seguins Pazzis's proof of Gerstenhaber's theorem [dSP13].

We also prove an induced correspondence between the properties of being acyclic and nil. Here, we need to reconcile the difference between conjugation (which preserves nilpotency) and congruence (where restrictions are natural), so the following alternative definition of induced subspaces is needed. This idea originates from the theory of non-commutative graphs [Wea21].

Definition 1.10. Let $\mathcal{S} \leq M(n, \mathbb{C})$ be a matrix space over the complex field \mathbb{C} . For a subspace $U \leq \mathbb{C}^n$ of dimension d , let T_U be an $d \times n$ matrix whose rows form an orthonormal basis of U . The *induced subspace* of \mathcal{S} on U is defined as $\mathcal{S}[U] := \{T_U B T_U^* \mid B \in \mathcal{S}\} \leq M(d, \mathbb{C})$.⁵

We discuss its differences with Definition 1.6 and why this is also a natural definition in Section 1.3. The induced correspondence between nil and acyclicity is the following, proved in Section 4.3.

Theorem 1.11. *Let $G = ([n], E)$ be a directed graph and $\mathcal{S}_G \leq M(n, \mathbb{C})$ be the graphical matrix space associated with G over \mathbb{C} . The maximum number of vertices in an acyclic induced subgraph of G equals the largest dimension of $U \leq \mathbb{C}^n$ such that $\mathcal{S}_G[U]$ is nil.*

1.2.3 Strong connectivity and irreducibility

A directed graph $G = ([n], E)$ is called *strongly connected* if for any non-empty and proper $S \subset [n]$, there are arcs going out of S . Equivalently, G is strongly connected if there is a directed path connecting any ordered pair of vertices. A matrix space $\mathcal{S} \leq M(n, \mathbb{F})$ is *irreducible*, if it has no non-trivial invariant subspace; otherwise we call it reducible. Namely, for any non-zero and proper subspace $U < \mathbb{F}^n$, $\mathcal{S}(U) := \langle \cup_{B \in \mathcal{S}} B(U) \rangle$ is not contained in U . Similar to the properties of being acyclic and nil, we prove basic, inherited, and induced correspondences between strong connectivity and irreducibility. We begin with the basic correspondence.

Theorem 1.12 (Simplified version of Theorem 5.1). *Let $G = ([n], E)$ be a directed graph and $\mathcal{S}_G \leq M(n, \mathbb{F})$ be the graphical matrix space associated with G . Then G is strongly connected if and only if \mathcal{S}_G is irreducible.*

Next, we state the inherited correspondence between strong connectivity and irreducibility, whose proof is given in Section 5.2.

Theorem 1.13. *Let $G = ([n], E)$ be a directed graph and $\mathcal{S}_G \leq M(n, \mathbb{F})$ be the graphical matrix space associated with G . The maximum number of arcs in a non-strongly connected subgraph of G equals the largest dimension of a reducible subspace of \mathcal{S}_G .*

⁵While different orthonormal bases lead to different subspaces of $M(d, \mathbb{C})$, these subspaces are unique up to conjugacy (and congruence) by invertible unitary matrices.

Finally, we state our induced correspondence result between strong connectivity and irreducibility, whose proof is in Section 5.3.

Theorem 1.14. *Let $G = ([n], E)$ be a directed graph and $\mathcal{S}_G \leq M(n, \mathbb{F})$ be the graphical matrix space associated with G . The maximum number of vertices in a non-strongly connected induced subgraph of G equals the largest dimension of $U \leq \mathbb{C}^n$ such that $\mathcal{S}_G[U]$ is reducible.*

Note that the minimum number of arcs that can be deleted from G to make it not strongly connected is a well-known graph-theoretic parameter called the *arc-strong connectivity* of G . Similarly, the minimum number of vertices whose deletion makes G not strongly connected is well-known as the *vertex-strong connectivity* of G . See e.g. [BJG08, Chapter 1.5]. In [LQ20], connections between the vertex and edge connectivities of undirected graphs and certain parameters of alternating matrix spaces were established. The proofs of Theorems 1.13 and 1.14 are non-trivial adaptations of those for [LQ20, Propositions 2.5 and 2.4].

1.2.4 Isomorphism, conjugacy, and congruence

Let $G = ([n], E)$ and $H = ([n], F)$ be two directed graphs. We say that G and H are *isomorphic*, if there exists a permutation $\sigma \in S_n$, such that $(i, j) \in E$ if and only if $(\sigma(i), \sigma(j)) \in F$.

For matrix spaces, there are two natural notions which can play the role of “isomorphism”. Let $\mathcal{S}_1, \mathcal{S}_2 \leq M(n, \mathbb{F})$ be two matrix spaces. We say that \mathcal{S}_1 and \mathcal{S}_2 are *conjugate*, if there exists $T \in GL(n, \mathbb{F})$, such that $\mathcal{S}_1 = T\mathcal{S}_2T^{-1} := \{TCT^{-1} \mid C \in \mathcal{S}_2\}$. We say that \mathcal{S}_1 and \mathcal{S}_2 are *congruent*, if there exists $T \in GL(n, \mathbb{F})$, such that $\mathcal{S}_1 = T\mathcal{S}_2T^t := \{TCT^t \mid C \in \mathcal{S}_2\}$.⁶ The following result establishes a basic correspondence between graph isomorphism and both notions of matrix space “isomorphism”.

Theorem 1.15 (Combined and simplified version of Propositions 6.2 and 6.9). *Let $G = ([n], E)$ and $H = ([n], F)$ be two directed graphs and $\mathcal{S}_G, \mathcal{S}_H \leq M(n, \mathbb{F})$ be their graphical matrix spaces, respectively. The following are equivalent:*

1. G and H are isomorphic.
2. \mathcal{S}_G and \mathcal{S}_H are conjugate.
3. \mathcal{S}_G and \mathcal{S}_H are congruent.

The proof of the equivalence between 1 and 3 in Theorem 1.15 is straightforward, and the proof strategy goes back to [HQ21]. We can actually prove a somewhat more general basic correspondence for congruence (see Proposition 6.2), which states that G is isomorphic to a subgraph of H if and only if \mathcal{S}_G is congruent to a subspace of \mathcal{S}_H .

The proof of the equivalence between 1 and 2 in Theorem 1.15 is elementary but much more complicated, and the proof strategy is inspired by [BS20, Theorem 4.13 in arXiv version 2]. Interestingly, the stronger correspondence mentioned in the previous paragraph fails in the conjugacy setting. That is, there exist graphs G and H such that H is not isomorphic to any subgraph of G but \mathcal{S}_H is conjugate to a subspace of \mathcal{S}_G (cf. Example 6.12). This suggests a subtle difference between these two symmetry notions of matrix spaces.

Interestingly, the basic correspondence between isomorphism and congruence *cannot* be boosted to an inherited correspondence, as seen in Example 6.5.

⁶When $\mathbb{F} = \mathbb{C}$, or more generally, \mathbb{F} is a quadratic extension of a subfield, we may adopt the conjugate transpose T^* instead of T^t .

1.2.5 Vertex transitivity and conjugacy/congruence irreducibility

Let $G = ([n], E)$ be a directed graph. We say that G is *vertex-transitive* if its automorphism group acts transitively on $[n]$. Recall that a matrix group $\mathcal{G} \leq \text{GL}(n, \mathbb{F})$ is reducible⁷, if there exists a non-zero and proper $U < \mathbb{F}^n$ such that for any $A \in \mathcal{G}$, $A(U) \leq U$. Otherwise, we call \mathcal{G} *irreducible*. Let $\mathcal{S} \leq \text{M}(n, \mathbb{F})$. Define $\text{Conj}(\mathcal{S}) := \{T \in \text{GL}(n, \mathbb{F}) \mid T\mathcal{S}T^{-1} = \mathcal{S}\} \leq \text{GL}(n, \mathbb{F})$. We say that \mathcal{S} is *conjugacy irreducible*, if $\text{Conj}(\mathcal{S})$ is irreducible as a matrix group. Define $\text{Cong}(\mathcal{S}) := \{T \in \text{GL}(n, \mathbb{F}) \mid T\mathcal{S}T^t = \mathcal{S}\} \leq \text{GL}(n, \mathbb{F})$. We say that \mathcal{S} is *congruence irreducible*, if $\text{Cong}(\mathcal{S})$ is irreducible as a matrix group. The following result establishes a basic correspondence between vertex transitivity and conjugacy/congruence irreducibility.

Theorem 1.16. *Let \mathbb{F} be a field of order > 2 , and let $G = ([n], E)$ be a directed graph. Then the following are equivalent:*

1. G is vertex-transitive.
2. \mathcal{S}_G is conjugacy irreducible.
3. \mathcal{S}_G is congruence irreducible.

See Section 7 for the proof. This generalizes the following result from quantum information. It is well-known that a directed graph can be embedded into a quantum channel. In [BBLM20], Bannink, Briët, Labib, and Maassen showed that a directed graph is vertex-transitive if and only if the resulting quantum channel is “irreducibly covariant”. Theorem 1.16 readily implies this result, and generalizes it from $\mathbb{F} = \mathbb{C}$ to \mathbb{F} being any field of order > 2 .

1.3 Perspective: symmetries of matrix spaces

In the results above, it was sometimes natural to view the restricted support of a matrix space as a bipartite graph, and sometimes as a directed graph. Moreover, the corresponding linear-algebraic properties were defined in a number of different ways. These choices stem from different natural ways of viewing the matrix space $\text{M}(n, \mathbb{F})$, which in turn are associated with different natural group actions on them, which may preserve some properties but not others. We now explain these in detail.

Let V and W be two n -dimensional vector spaces over a field \mathbb{F} . Then there are (at least) three natural ways of viewing the matrix space $\text{M}(n, \mathbb{F})$. Associated to each of these interpretations is a natural group action, which may not preserve every natural property of matrices, but does preserve all the natural properties which are inherent to the interpretation.

- First, we may identify $\text{M}(n, \mathbb{F})$ with the set of linear maps $V \rightarrow W$. In this case, the natural group action on $\text{M}(n, \mathbb{F})$ is the left-right action of the group $\text{GL}(n, \mathbb{F}) \times \text{GL}(n, \mathbb{F})$, which corresponds to changing the bases of V and W . This group action preserves natural properties of linear maps $V \rightarrow W$, such as their rank.
- Secondly, we may identify $\text{M}(n, \mathbb{F})$ with the set $\text{End}(V)$ of endomorphisms of V , i.e., linear maps from V to itself. In this case, the natural group action is the conjugation action of $\text{GL}(n, \mathbb{F})$, which corresponds to changing the basis of V . This symmetry preserves additional properties, such as the spectrum, determinant, and nilpotency.

⁷Note that this definition agrees with that of reducibility in Section 1.2.3: $\mathcal{G} \leq \text{GL}(n, \mathbb{F})$ is reducible if and only if the linear subspace $\langle \mathcal{G} \rangle \leq \text{M}(n, \mathbb{F})$ it spans is reducible as a matrix space.

- Finally, we may identify $M(n, \mathbb{F})$ with the set $\text{Bil}(V)$ of bilinear forms $V \times V \rightarrow \mathbb{F}$. In this case, the natural action is the congruence action of $\text{GL}(n, \mathbb{F})$, which arises from changing the basis of the Gram matrices associated with the bilinear forms. This symmetry preserves natural properties of bilinear forms, such as symmetry and non-degeneracy, but does not preserve other properties such as nilpotency.

Similarly, the natural group action on n -vertex directed graphs is that of S_n (permuting the vertices), and the natural group action on bipartite graphs with m vertices on the left and n vertices on the right in each part is that of $S_m \times S_n$. In each of our correspondences between graph theory and linear algebra, there is an underlying interpretation of $M(n, \mathbb{F})$, and correspondingly an underlying relationship between the symmetries. For example, in Theorem 1.5, we view $M(m \times n, \mathbb{F})$ as the space of linear maps from an m -dimensional vector space to an n -dimensional vector space, with the symmetries being the left-right action of $\text{GL}(m, \mathbb{F}) \times \text{GL}(n, \mathbb{F})$. Here, the (bipartite) graph property of matching size is preserved by the action of $S_m \times S_n$, and the corresponding matrix property of rank is preserved by the left-right action.

In contrast, in Theorem 1.8, we view $M(n, \mathbb{F})$ as the space of endomorphisms of an n -dimensional vector space, with the conjugation action of $\text{GL}(n, \mathbb{F})$. This symmetry preserves the property of nilpotency, just as the corresponding action of S_n on the set of directed graphs preserves the property of being acyclic.

Finally, in the correspondence between directed graph isomorphism and matrix space congruence (Theorem 1.15), we use the third interpretation of $M(n, \mathbb{F})$, namely the identification with $\text{Bil}(V)$. This is immediately apparent from the fact that the definition of congruence concerns the natural symmetries of this interpretation.

Note that the third interpretation naturally comes with a notion of restriction. Indeed, let V and W be vector spaces over \mathbb{F} , and $U_1 \leq V$ and $U_2 \leq W$ be two subspaces. Then a bilinear form $V \times W \rightarrow \mathbb{F}$ immediately restricts to a bilinear form $U_1 \times U_2 \rightarrow \mathbb{F}$. This is why Definition 1.6 is natural for bilinear forms. In contrast, for a linear map $f : V \rightarrow V$, one cannot restrict the codomain from V to $U \leq V$ unless U is an invariant subspace of V under f .

This is relevant for our two induced correspondences, Theorems 1.11 and 1.14. In both of them, we study properties of endomorphisms (nilpotency and irreducibility, respectively), and thus the symmetry must preserve these properties, which means that the underlying symmetry is the conjugation action of $\text{GL}(n, \mathbb{F})$. However, because they establish a correspondence with induced subgraphs, we also need a notion of restriction. But as discussed above, in general, restriction does not preserve these properties of endomorphisms, as the natural symmetry is that of $\text{Bil}(V)$, rather than $\text{End}(V)$. Thus, for these induced correspondences, we need a symmetry that is simultaneously a conjugation action and a congruence action. The natural such symmetry is that of the unitary group. This is why induced subspaces (Definition 1.10) are only defined when the underlying field is \mathbb{C} , and why Theorems 1.11 and 1.14 are only stated over \mathbb{C} as well: it is only when we have access to a symmetry that preserves *both* the structure of $\text{End}(V)$ and $\text{Bil}(V)$ that we can expect such induced correspondences to hold. Moreover, this explains why Definition 1.10 is natural: by insisting that the rows of T_U form an orthonormal basis of U , we are ensuring that the operation of passing to an induced subspace preserves all the relevant properties.

As a final remark, we think it is interesting and noteworthy that in the first two interpretations of matrix spaces, we have examples of inherited correspondences (the rank-matching correspondence for the first interpretation, and the acyclicity-nilpotency and strong connectivity-irreducibility correspondences for the second interpretation). In contrast, under the third interpretation, we have an instance of a basic correspondence which *cannot* be boosted to an inherited correspondence (isomorphism and congruence). It would be extremely interesting to classify which basic corre-

spondences yield inherited correspondences and which ones do not; it seems possible that the first two interpretations are more amenable to such inherited correspondences.

1.4 Perspective: computational complexity

Works based on the connection between matchings and ranks. In Section 1.1, we reviewed how the connection between perfect matchings and full-rank matrices served as a key to several important questions and results in computational complexity. Lovász [Lov79] used this correspondence to develop an RNC algorithm for the bipartite matching problem. The question of computing a perfect matching (if there exists one) in RNC was subsequently solved by Karp, Upfal and Wigderson [KUW86], then simplified by Mulmuley, Vazirani and Vazirani [MVV87]. Recently, deterministic quasi-NC algorithms were devised for the perfect matching problem on bipartite graphs [FGS19] and on general graphs [ST17].

Based on the connection between bipartite matching and singularity, Edmonds [Edm67] proposed the question of finding an efficient deterministic algorithm for determining if a matrix space contains a non-singular matrix (or, equivalently, if a symbolic matrix has a non-zero determinant). Over finite fields this problem is NP-complete [BFS99], and over large enough fields this problem is in RP [Lov79] by the Schwartz–Zippel lemma [Sch80, Zip79]. The converse problem, namely to decide if a matrix space contains only singular matrices, is now more commonly known as the symbolic determinant identity testing (SDIT) problem. To derandomize SDIT is a central problem in computational complexity. It is equivalent to the polynomial identity testing problem for algebraic branching programs [SY09]. Its central importance in computational complexity stems from the following surprising result of Kabanets and Impagliazzo [KI04]: An efficient deterministic algorithm for SDIT implies strong circuit lower bounds which seem beyond the reach of current techniques. Underlying this is a foundational result of Valiant [Val79], who found a way of converting a short algebraic formula for a polynomial f into a small symbolic matrix B_f with $\det(B_f) = f$. An important aspect of Valiant’s construction is that the symbolic matrix B_f has a very restricted support; in fact, in his original paper, he explicitly viewed the entries of B_f as living on the edges of a graph! This yields another reason to study matrix spaces with restricted support: They may be easier to understand, and nonetheless a deterministic algorithm for SDIT with certain restrictions on the support would yield a deterministic algorithm for polynomial identity testing. Additionally, it is possible that finding an algorithm for SDIT is easier in certain classes of restricted support, and studying such intermediate questions is a natural approach towards the general problem. One example of restricted support that is natural to study is that which captures *graph rigidity*, itself a celebrated open problem (see e.g. [RW19]).

The non-commutative version of the SDIT problem, known as the non-commutative rank problem of symbolic matrices, has received considerable attention recently. It can be viewed as a linear algebraic analog of the problem of deciding if a bipartite graph has “Hall’s obstructions” [Hal35] for perfect matchings. Originally proposed by P. M. Cohn [Coh75] in the context of free skew fields, this problem was recently shown to be in P via three solutions: over \mathbb{Q} by [GGOW20], and over any field first by [IQS17, IQS18] and then by [HH21].

Discussions on the connection between cyclicity and nilpotency. After identifying the correspondence between cycles and nilpotent matrices, it is natural to examine the following questions.

1. Compute a cycle in a directed graph in NC.
2. Decide if a matrix space is nil, i.e., consisting of only nilpotent matrices. We call this the

Symbolic Nil Testing (SNT) problem

It turns out that the first problem reduces to computing a perfect matching on a bipartite graph in NC (Proposition A.1), and the second problem reduces to the symbolic determinant identity testing (SDIT) problem (Proposition A.3). In fact, it can be shown that computing a cycle in a directed graph can be done in NC (Proposition A.2).

It is interesting to compare SNT and SDIT from the perspective of deterministic algorithms. For example, singular matrix spaces are preserved under the left-right action, and the nullcone of the left-right action on matrix tuples consists of those with a “shrunk subspace” (a linear-algebraic analog of the “Hall’s obstruction” for perfect matching). To determine if a matrix tuple lies in the nullcone of the left-right action is just the non-commutative rank problem, which is a rather non-trivial problem and we briefly touched upon this in the above. To solve this problem in a black-box way is still open.

On the other hand, nil matrix spaces are preserved under the conjugation action, and the nullcone of the conjugation action on matrix tuples consists of those which can be simultaneously strict upper-triangularized. To determine if a matrix tuple $\mathbf{B} = (B_1, \dots, B_m) \in M(n, \mathbb{F})^m$ can be simultaneously strict upper-triangularized can be solved by testing if $\mathbf{B}^n = \langle B_{i_1} B_{i_2} \cdots B_{i_n} \mid i_j \in [m] \rangle$ is the zero matrix space. A black-box quasipolynomial-time algorithm for the Noether normalization of this action was shown by Forbes and Shpilka [FS13].

As another example, consider those matrix spaces that have a basis of rank-1 matrices. In this case, SDIT can be solved in deterministic polynomial time, again with non-trivial algorithms [Gur04, IKS10, IKQS15]. For SNT, this is easy because it is also strictly upper-triangularizable [MOR91], so the algorithm mentioned in the above paragraph suffices.

To summarize, SNT seems easier than SDIT for most special cases. Given these, it may be surprising to note that SDIT reduces to a quantitative version of SNT as follows. Recall that a symbolic matrix is a matrix whose entries are affine linear forms. Let B be a symbolic matrix. SDIT asks if $\det(B)$ is the zero polynomial. The nilpotency index of B is the smallest k such that B^k is the zero matrix.

Proposition 1.17. *SDIT reduces to the problem of deciding if the nilpotency index of a symbolic matrix is no more than a given integer.*

Proof. Let B be a symbolic matrix of size $n \times n$. By [Ber84], there exist symbolic matrices C_1, \dots, C_ℓ of size t , $\ell, t = \text{poly}(n)$, such that the $(1, 1)$ entry of $C_1 \cdots C_\ell$ is $\det(B)$. Let $C_0 = C_{\ell+1} = E_{1,1}$. Construct a $(\ell + 3) \times (\ell + 3)$ block matrix T with the block sizes being t . For $i \in [\ell + 2]$, the $(i, i + 1)$ th block of T is C_{i-1} . We then see that the $(1, \ell + 3)$ th block of $T^{\ell+2}$ is $C_0 \cdot C_1 \cdots C_{\ell+1}$, which is zero if and only if $\det(B)$ is the zero polynomial. So the nilpotency index of T is $\leq \ell + 2$ if and only if $\det(B)$ is the zero polynomial. This concludes the proof. \square

Note that the symbolic matrix constructed in the proof of Proposition 1.17 is clearly nilpotent. Therefore, it does not give us a reduction from SDIT to SNT.

From the above discussion, at present it looks to us that solving the general SNT is a good challenge. This is because classifying nil matrix spaces is a wild problem [Ger58] with implications [VF17] to the classical Albert’s problem [Alb50]. This is also motivated by the work of Makam and Wigderson [MW21], who showed that the matrix tuples spanning singular matrix spaces do not form a nullcone in general. Therefore, it is interesting to identify further algebraic varieties for which the membership problem is interesting enough as the test bed for further progress towards solving SDIT.

Applications: NP-hardness for some linear-algebraic problems. One consequence of our inherited and induced correspondences is that we can prove NP-hardness results for several linear-algebraic problems, using known NP-hardness results for certain graph-theoretic problems. Namely, we have the following results.

Theorem 1.18. *Let $\mathcal{S} \leq M(n, \mathbb{F})$ be of dimension m whose matrices have rank at most r . Let $d \in [m]$. To decide if \mathcal{S} admits a dimension- d subspace whose matrices have rank at most $r - 1$ is NP-hard over any field, and NP-complete over finite fields.*

Proof. For a bipartite graph $G = (L \cup R, E)$, a k -blocker is a subgraph of G whose matching number is at most the matching number of G minus k . It was shown in [ZRP⁺09, Theorem 3.3] that deciding the existence of a size- d 1-blocker in bipartite graphs is NP-complete. The theorem then follows with Theorem 1.5. \square

Theorem 1.19. *Let $\mathcal{S} \leq M(n, \mathbb{F})$ be of dimension m , and $d \in [m]$. To decide if \mathcal{S} admits a dimension- d nil subspace is NP-hard over any field, and NP-complete over finite fields.*

Proof. This follows from Theorem 1.9 and Karp’s classical result that the feedback arc set problem is NP-complete [Kar72]. \square

Theorem 1.20. *Let $\mathcal{S} \leq M(n, \mathbb{C})$, and $d \in [n]$. To decide if \mathcal{S} admits a d -dimensional $U \leq \mathbb{C}^n$ such that $\mathcal{S}[U]$ is nil is NP-hard.*

Proof. This follows from Theorem 1.9 and Karp’s well-known result that the feedback vertex set problem is NP-complete [Kar72]. \square

Theorem 1.21. *Let $\mathcal{S}_1, \mathcal{S}_2 \leq M(n, \mathbb{F})$. To decide if \mathcal{S}_1 contains a congruent copy of \mathcal{S}_2 as a subspace is NP-hard over any field, and NP-complete over finite fields.*

Proof. This follows from the NP-completeness of the subgraph isomorphism problem (via e.g. Hamiltonian cycles or paths [Kar72]) and Proposition 6.2 (mentioned after Theorem 1.15). \square

1.5 Connections with quantum information theory

For this subsection, we restrict to the case $\mathbb{F} = \mathbb{C}$. Graph theory has inspired several research topics in quantum information. For example, Hastings [Has07], and Ben-Aroya, Schwartz, and Ta-Shma [BASTS10] introduced quantum expanders which are quantum channels satisfying certain regular and expanding properties. Roughly speaking, they view quantum channels as generalizations of the adjacency matrices of graphs. This leads to other research lines linking quantum information theory with graph theory, which indicate that some spectral properties of quantum channels properly generalize certain graph-theoretic properties (of adjacency matrices).

As another example, in the study of zero-error communications of quantum channels, Duan, Severini and Winter in [DSW13] associated each quantum channel an operator system, which is analog to the classical setting where Shannon associated each classical channel with a graph [Sha56]. This gives rise to another way to generalize graphs to matrix spaces. Along with this generalization, Duan, Severini, and Winter proposed quantum generalizations of independence number and the Lovász theta function for operator systems. Viewing operator systems as generalizations of graphs has led to a fruitful research line which connects combinatorics, operator algebra and optimization theory; see [Wea21] and the references therein.

We point out that the results presented in this paper recover and generalize several interesting results relating to the above two types of generalizations. Thus, we believe that our framework

provides a more systematical way to study connections between graphs and matrix space. To be more specific, we show that Theorems 1.12, 1.15 and 1.16 can be applied to the quantum settings, which reveals that

- the irreducibility of quantum channels (introduced in [EHK78, Wol12]) is a generalization of the strong connectivity of directed graphs (Corollary 8.2);
- the irreducibly covariant quantum channels are generalizations of vertex transitive graphs (Corollary 8.4, first proved in [BBLM20]);
- the connectivity of operator systems (introduced in [CDS21]) is a generalization of connectivity of (undirected) graphs (Corollary 8.7, first proved in [CDS21]);
- the isomorphism between operator systems (introduced in [OP15]) is a generalization of the isomorphism between graphs (Corollary 8.9, first proved in [OP15]).

We emphasize that all the above generalizations are already shown in the references. However, our results simplify some of the proofs and obtain more general results. See Section 8 for a more detailed discussion.

1.6 Some known connections between graphs and matrix spaces

Inspired by the classical connection between perfect matchings and full-rank matrices, several correspondences between graph-theoretic structures, and structures for alternating matrix spaces, have been discovered recently, including:

- Independent sets vs. (totally) isotropic spaces, and vertex colorings vs. (totally) isotropic decompositions [BCG⁺21].
- Connectivity vs. orthogonal indecomposability. As a consequence, correspondences of vertex and edge connectivities for alternating matrix spaces are also presented [LQ20].
- Isomorphism notions for graphs and alternating matrix spaces [HQ21].

Some graph-theoretic questions were also translated to the matrix space setting, including:

- Transferring techniques for graph isomorphism to study matrix space congruence [LQ17].
- Enumeration formulas of isotropic spaces and orthogonal decompositions as q -analogs of enumeration formulas of independent sets and connected graphs [Qia21].
- Turán and Ramsey problems for alternating matrix spaces [Qia20].
- New graph polynomials from group zeta functions, through the connection between alternating matrix spaces and nilpotent p -groups [RV22, Ros22].
- A geometric characterization of graph k -connectivity, which can be equivalently phrased as a linear-algebraic property of the matrix space associated to the graph [LLW88].
- Connections between graph-theoretic and linear-algebraic notions of expansion [BBLM20, LQW⁺22]; see Section 8.2 for more details.

Some results and techniques in the above works inspire results in the present work. One key difference is that one emphasis in this paper is to examine the connection between *directed graphs* and matrix spaces, whereas most works in the above-studied connections between undirected graphs and bipartite graphs and matrix spaces.

1.7 Discussion, research directions, and open problems

As we have attempted to demonstrate in the introduction, there is a surprising wealth of correspondences between graph theory and matrix spaces. Our goal in this paper is to initiate a systematic study of such correspondences, and we hope that this will lead to fruitful future work. One big “meta-question” that we are interested in is a further development of this theory. Which other natural graph-theoretic properties have linear-algebraic analogs? Conversely, which linear-algebraic properties correspond to graph-theoretic ones in matrix spaces of restricted support? More broadly, can one develop a general theory and characterize the family of such properties? Additionally, is it possible to characterize which basic correspondences can be boosted to inherited and/or induced correspondences? Given that two of our inherited correspondences (Theorems 1.5 and 1.8) generalize classical results on matrix spaces (Dieudonné’s theorem [Die48] and Gerstenhaber’s theorem [Ger58], respectively), it is possible that such a general theory could imply or generalize important results in algebra.

It would also be interesting to find further applications of such correspondences. As mentioned above, the basic correspondence between bipartite perfect matchings and singularity of matrix spaces has been immensely fruitful, yielding in particular an efficient randomized parallel algorithm for perfect matching [Lov79]. As discussed in Section 1.4, our new results allow us to prove NP-hardness results for certain linear-algebraic questions, and lead to the matrix space nil test problem which is rather interesting in the context of polynomial identity testing. In Section 1.5, we explained how our results generalize certain results in quantum information theory. The reason for applications in quantum information is that matrix tuples and matrix spaces are used to set up superoperators there. Since matrix spaces arise in numerous mathematical areas as indicated at the beginning of Section 1.1, with the number and depth of the correspondences that are now known to exist, it seems likely that there are further applications waiting to be discovered.

Finally, we mention one specific conjecture whose proof or disproof we would be extremely interested in. Recall that Theorem 1.8 gives a basic correspondence between the properties of being acyclic and nil. There is a simple natural generalization of this, which we prove in Section 4.1.2, which says the following for a directed graph $G = ([n], E)$ and any non-negative integer k . Every collection of disjoint cycles in G covers at most k vertices if and only if every matrix in \mathcal{S}_G has at most k non-zero eigenvalues. Note that the case $k = 0$ corresponds precisely to G being acyclic, and to \mathcal{S}_G being nil. Our conjecture is that this basic correspondence can be boosted to an inherited correspondence.

Conjecture 1.22. *Let $G = ([n], E)$ be a directed graph, and let $0 \leq k \leq n$ be an integer. The maximum number of edges in a subgraph H of G in which every collection of disjoint cycles covers at most k vertices equals the largest dimension of a subspace of \mathcal{S}_G in which every matrix has at most k non-zero eigenvalues.*

If true, Conjecture 1.22 would generalize a well-known theorem of Atkinson [Atk80], who proved that if \mathbb{F} is sufficiently large and $S \leq M(n, \mathbb{F})$ is a matrix space in which every matrix has at most k non-zero eigenvalues, then $\dim S \leq nk + \binom{n-k}{2}$. This is precisely the statement of Conjecture 1.22 in case G is a complete directed graph, i.e., Conjecture 1.22 is the restricted-support version of Atkinson’s full-support theorem. Note too that Conjecture 1.22 generalizes Theorem 1.8, just as Atkinson’s theorem generalizes Gerstenhaber’s theorem.

Some algorithmic problems about matrix spaces are also worth studying. For example, Theorem 1.13 naturally leads to the problem of finding the maximum dimension over reducible subspaces of a matrix space. The corresponding graph-theoretic problem, namely the min-cut problem for

directed graphs, can be solved in deterministic polynomial time [BJG08]. The submodular optimization algorithms over modular lattices [HH21] may be relevant for this purpose.

2 Preliminaries

We collect basic notation. More definitions and notation will be introduced in each section.

For $n \in \mathbb{N}$, $[n] := \{1, 2, \dots, n\}$.

Graphs. A directed graph is $G = (V, E)$, where V is the vertex set and $E \subseteq V \times V$ is the *arc* set. We shall mostly working with directed graphs whose vertex set is $V = [n]$.

A bipartite graph is $G = (L \cup R, E)$, where L and R are the left and right vertex sets, and $E \subseteq L \times R$ is the *edge* set. We shall mostly working with bipartite graphs with $L = R = [n]$.

An undirected graph is $G = (V, E)$, where V is the vertex set and $E \subseteq \binom{V}{2} := \{\{v, v'\} \mid v, v' \in V, v \neq v'\}$ is the *edge* set.

We say two graphs $G = ([n], E)$ and $H = ([n], F)$ are isomorphic if there is a bijection $f : [n] \rightarrow [n]$ such that $(i, j) \in E \Leftrightarrow (f(i), f(j)) \in F$.

We denote E' being a subset of E as $E' \subseteq E$ and E' being a *proper* subset of E as $E' \subset E$. A spanning subgraph $G' = (V, E')$ of $G = (V, E)$ has the same vertex set V and edge set $E' \subseteq E$. Sometimes we call spanning subgraph as subgraph if there is no confusion. An induced subgraph $G[V'] = (V', E')$ of $G = (V, E)$ has the vertex set $V' \subseteq V$, and $E' = \{(v, v') \in E \mid v, v' \in V'\}$.

Vector spaces. Let \mathbb{F} be a field. Let \mathbb{F}^n be the vector space of length- n *row* vectors. The i th standard basis vector of \mathbb{F}^n is denoted e_i . We use $U \leq \mathbb{F}^n$ to denote that U is a subspace of \mathbb{F}^n (and use $<$ for *proper* subspace). For a set of vectors $\{v_1, \dots, v_m\}$, let $\langle v_1, \dots, v_m \rangle \leq \mathbb{F}^n$ be their linear span.

Matrices. Let $M(n, \mathbb{F})$ be the linear space of $n \times n$ matrices over \mathbb{F} . Given $(i, j) \in [n] \times [n]$, the elementary matrix $E_{i,j} \in M(n, \mathbb{F})$ is the matrix with the (i, j) th entry being 1, and other entries 0. We use I_n to denote the $n \times n$ identity matrix.

For a matrix $B \in M(n, \mathbb{F})$, let $\text{rowspan}(B)$ be the vector space spanned by the rows of B and $\text{colspan}(B)$ be the vector space spanned by the columns of B . We have $\text{rank}(B) = \dim(\text{colspan}(B)) = \dim(\text{rowspan}(B))$. A matrix $B \in M(n, \mathbb{F})$ is nilpotent, if $B^k = 0$ for some $k \in \mathbb{N}$. For $i, j \in [n]$, $B(i, j)$ denotes the (i, j) th entry of B .

For a matrix $B \in M(n, \mathbb{F})$ and a vector $v \in \mathbb{F}^n$, we usually consider the right action of B on v , and denote the result as vB . The reason to use row vectors and matrices acting on the right is to be consistent with directed graphs: if we use $E_{i,j}$ to represent the arc from i to j , then the right action of $E_{i,j}$ sends e_i to e_j . For a subspace $U \leq \mathbb{F}^n$, let $B(U) := \{uB \mid u \in U\} \leq \mathbb{F}^n$.

Matrix tuples and matrix spaces. A matrix tuple \mathbf{S} of length m is an element of $M(n, \mathbb{F})^m$. A *matrix space* \mathcal{S} is a linear subspace of $M(n, \mathbb{F})$. Given $\mathcal{S} \leq M(n, \mathbb{F})$ and $R, T \in M(n, \mathbb{F})$, $R\mathcal{S}T := \{RST \mid S \in \mathcal{S}\} \leq M(n, \mathbb{F})$. Two matrix spaces $\mathcal{S}, \hat{\mathcal{S}} \leq M(n, \mathbb{F})$ are *conjugate* if there exists $T \in \text{GL}(n, \mathbb{F})$ such that $T\mathcal{S} = \hat{\mathcal{S}}T$. They are *congruent* if there exists $T \in \text{GL}(n, \mathbb{F})$ such that $T\mathcal{S}T^t = \hat{\mathcal{S}}$.

3 Matchings in bipartite graphs and ranks of matrices

3.1 The inherited correspondence

In this subsection, we prove Theorem 1.5, establishing an inherited correspondence between the matching size in a bipartite graph and the rank of matrices in its graphical matrix space.

Let $G = ([m] \cup [n], E)$ be a bipartite graph with $m \leq n$.

A matching in G is defined as a subset $M \subseteq E$ such that for any $(i_1, j_1), (i_2, j_2) \in M$, $i_1 \neq i_2$ and $j_1 \neq j_2$.

Let $B \in M(m \times n, \mathbb{F})$ be the bipartite adjacency matrix of $G = ([m] \cup [n], E)$, where $B(i, j) = 1$ if and only if $(i, j) \in E$; and $B(i, j) = 0$ otherwise. Let $\rho(B)$ be the minimum number of columns and rows which cover all non-zero entries of B . Then König's theorem, a fundamental result in graph theory, says that the matching number of G equals $\rho(B)$ (e.g. see [Die17, Theorem 2.1.1]).

For $(i_1, j_1), (i_2, j_2) \in [m] \times [n]$, the lexicographic order \prec is defined as $(i_1, j_1) \prec (i_2, j_2)$ if and only if $i_1 < i_2$ or $i_1 = i_2$ and $j_1 < j_2$. For $A \in M(m \times n, \mathbb{F})$, let $p(A) = \min\{(i, j) : A(i, j) \neq 0\}$ be the position of A 's lexicographically first non-zero entry. An important result in [Mes85] is the following:

Lemma 3.1 ([Mes85, Theorem 1]). *Given matrices $A_1, \dots, A_d \in M(m \times n, \mathbb{F})$, construct $B \in M(m \times n, \mathbb{F})$ to have 1 in all the positions $p(A_1), \dots, p(A_d)$ and zeros elsewhere. Then $\langle A_1, \dots, A_d \rangle$ contains a matrix of rank at least $\rho(B)$.*

Let $\text{MaxBdMatSize}(G, r)$ be the maximum size of a spanning subgraph of G whose matching number is at most r and $\text{MaxBdRankDim}(\mathcal{S}_G, r)$ be the largest dimension of a subspace of \mathcal{S}_G in which any matrix has rank at most r . We are ready to establish the following result.

Theorem 1.5, restated. *Let $G = ([m] \cup [n], E)$ be a bipartite graph with $m \leq n$, and $\mathcal{S}_G \leq M(m \times n, \mathbb{F})$ be the graphical matrix space associated with G . Then for each $r \in [m]$, $\text{MaxBdMatSize}(G, r) = \text{MaxBdRankDim}(\mathcal{S}_G, r)$.*

Proof. To see $\text{MaxBdMatSize}(G, r) \leq \text{MaxBdRankDim}(\mathcal{S}_G, r)$, take a maximum spanning subgraph G' with matching number $\leq r$. By Theorem 1.4, we know that any matrix in $\mathcal{S}_{G'}$ has rank $\leq r$. Thus, $\text{MaxBdMatSize}(G, r) = \dim(\mathcal{S}_{G'}) \leq \text{MaxBdRankDim}(\mathcal{S}_G, r)$.

To see $\text{MaxBdRankDim}(\mathcal{S}_G, r) \leq \text{MaxBdMatSize}(G, r)$, take the largest subspace W of \mathcal{S}_G in which any matrix has rank $\leq r$. Let $d = \dim(W)$. Without loss of generality, we choose a basis A_1, \dots, A_d of W such that $p(A_1), \dots, p(A_d)$ are all distinct. This can be done by Gaussian elimination (and viewing A_1, \dots, A_d as mn -dimensional vectors). Construct B as in Lemma 3.1 and let H be the bipartite graph whose bipartite adjacency matrix is B . By Lemma 3.1, $W = \langle A_1, \dots, A_d \rangle$ contains a matrix of rank at least $\rho(B)$, which is the matching number of H . Thus, the matching number of H cannot be greater than r , otherwise W would contain a matrix of rank greater than r . Since B has exactly d non-zero entries, H has exactly d edges. Moreover, since each A_i was supported on the edges of G , B is supported on the edges of G as well, and thus H is a spanning subgraph of G . These imply that $\text{MaxBdRankDim}(\mathcal{S}_G) \leq \text{MaxBdMatSize}(G)$. \square

Remark 3.2. Theorem 1.5 can be seen as a generalization of the Dieudonné–Flanders–Meshulam theorem [Die48, Fla62, Mes85], which states that the largest dimension of a matrix space of $M(m \times n, \mathbb{F})$ ($m \leq n$) with bounded rank r is rn . Taking G to be the complete bipartite graph, then $\mathcal{S}_G = M(m \times n, \mathbb{F})$. In this case, $\text{MaxBdRankDim}(\mathcal{S}_G, r)$ is given by the maximum size of spanning subgraphs of G whose matching number is r , which is exactly rn . This recovers the Dieudonné–Flanders–Meshulam theorem.

3.2 The induced correspondence

In this subsection, we prove Theorem 1.7, establishing an induced correspondence between the matching size in a bipartite graph and the rank of matrices in its graphical matrix space.

Suppose we have $\mathcal{S} \leq M(m \times n, \mathbb{F})$, $L \leq \mathbb{F}^m$, $R \leq \mathbb{F}^n$, $\dim(L) = s$ and $\dim(R) = t$. The *order* of \mathcal{S} is $m + n$. Let T_L (resp. T_R) be an $s \times m$ (resp. $t \times n$) matrix whose rows span L (resp. R). Define $\mathcal{S}[L, R] := \{T_L B T_R^t \mid B \in \mathcal{S}\}$ be the induced subspace of \mathcal{S} with respect to L and R .

Let $r \in \mathbb{N}$. For a graph G , let $\text{MaxBdMatOrd}(G, r)$ be the maximum order over induced subgraphs of G with matching number at most r . For a matrix space \mathcal{S} , let $\text{MaxBdRankOrd}(\mathcal{S}, r)$ be the maximum order over induced subspaces of \mathcal{S} with maximum rank at most r .

Theorem 1.7, restated. *Let $G = ([m] \cup [n], E)$ be a bipartite graph with $m \leq n$, and $\mathcal{S}_G \leq M(m \times n, \mathbb{F})$ be the graphical matrix space associated with G . Then for each $r \in [m]$, $\text{MaxBdMatOrd}(G, r) = \text{MaxBdRankOrd}(\mathcal{S}_G, r)$.*

Proof. To see $\text{MaxBdMatOrd}(G, r) \leq \text{MaxBdRankOrd}(\mathcal{S}_G, r)$, let $V_1 \cup V_2$ where $V_1 \subseteq [m]$ and $V_2 \subseteq [n]$ be vertex subsets of maximum order such that the matching number of $G[V_1 \cup V_2]$ is at most r . Take $L = \langle e_i \mid i \in V_1 \rangle \leq \mathbb{F}^m$ and $R = \langle e_i \mid i \in V_2 \rangle \leq \mathbb{F}^n$. Also take an $s \times m$ matrix T_L whose rows are exactly $\{e_i \mid i \in V_1\}$ and a $t \times n$ matrix T_R whose rows are exactly $\{e_i \mid i \in V_2\}$. By Theorem 1.4, any matrix in $\mathcal{S}_G[L, R]$ is of rank at most r , as $\mathcal{S}_G[L, R] = \mathcal{S}_{G[V_1 \cup V_2]}$ and any matching of $G[V_1 \cup V_2]$ has size at most r .

To see $\text{MaxBdRankOrd}(\mathcal{S}_G, r) \leq \text{MaxBdMatOrd}(G, r)$, let $L \leq \mathbb{F}^m$ with $\dim(L) = s$, $R \leq \mathbb{F}^n$ with $\dim(R) = t$, and $\dim(L) + \dim(R) = s + t = \text{MaxBdRankOrd}(\mathcal{S}_G, r)$ such that $\mathcal{S}_G[L, R]$ is an induced subspace of \mathcal{S}_G in which any matrix has rank at most r . Let T_L be an $s \times m$ matrix whose rows span L and T_R be a $t \times n$ matrix whose rows span R . Then there exists $V_1 \subseteq [m]$ (resp. $V_2 \subseteq [n]$) such that the columns in T_L (resp. T_R) with indices in V_1 (resp. V_2) are linearly independent. We want to show that the matching number of $G[V_1 \cup V_2]$ is at most r . Let A (resp. B) be the $s \times s$ (resp. $t \times t$) submatrix of T_L (resp. T_R) with column indices in V_1 (resp. V_2). This ensures that A and B are both full-rank.

Now we claim that $A \mathcal{S}_{G[V_1 \cup V_2]} B^t$ is a subspace of $\mathcal{S}_G[L, R]$. To see this, we denote by E_L (resp. E_R) the $m \times s$ (resp. $n \times t$) matrix formed of columns $\{e_i^t : i \in V_1\}$ (resp. $\{e_i^t : i \in V_2\}$) such that $A = T_L E_L$ (resp. $B^t = E_R^t T_R^t$). Note that $\mathcal{S}_{G[V_1 \cup V_2]} = E_L^t \mathcal{S}_G E_R$ and $A \mathcal{S}_{G[V_1 \cup V_2]} B^t = T_L E_L E_L^t \mathcal{S}_G E_R E_R^t T_R^t$. Since $E_L E_L^t$ and $E_R E_R^t$ are both diagonal, we have that $E_L E_L^t \mathcal{S}_G E_R E_R^t \leq \mathcal{S}_G$, implying $A \mathcal{S}_{G[V_1 \cup V_2]} B^t \leq T_L \mathcal{S}_G T_R^t = \mathcal{S}_G[L, R]$. It follows that the maximal rank of matrices in $\mathcal{S}_{G[V_1 \cup V_2]}$ is at most the maximal rank of matrices in $\mathcal{S}_G[L, R]$, which is at most r . By Theorem 1.4, $G[V_1 \cup V_2]$ is an induced subgraph of order- $(s + t)$ whose matching number is at most r , which concludes the proof. \square

4 Cycles in directed graphs and nilpotent matrices

In this section, we prove correspondences between cycles in directed graphs and nilpotency of matrices in their associated graphical matrix spaces. We begin with the basic correspondences, then prove the inherited correspondence in Section 4.2 and the induced correspondence in Section 4.3.

4.1 The basic correspondences

4.1.1 Maximum walk lengths and nil/nilpotent indices

Let $G = ([n], E)$ be a directed graph. A *walk of length k* in G is a sequence of $k + 1$ vertices $v_0, v_1, \dots, v_k \in [n]$, where $(v_{i-1}, v_i) \in E$ for all $i \in [k]$. We call v_0 and v_k the starting and ending vertices, respectively, and v_1, \dots, v_{k-1} are called the intermediate vertices. A *path* in G is a walk with no repeated vertices. A *cycle* in G is a path v_0, \dots, v_k with $v_0 = v_k$. We say that G is *cyclic* if there exists a cycle in G , and *acyclic* otherwise.

Let $\text{MaxWalkLen}(G)$ be the maximum length of walks in a directed graph $G = ([n], E)$. If G is cyclic, define $\text{MaxWalkLen}(G) = \infty$. From a walk of length n in G we can obtain a cycle in G . Thus for an acyclic graph G , $\text{MaxWalkLen}(G) \leq n - 1$. Moreover, if G is acyclic, every walk in G is a path. We shall relate cycles in directed graphs with nilpotent matrices in matrix spaces.

Definition 4.1. A matrix space $\mathcal{S} \leq M(n, \mathbb{F})$ is *nil*, if any $B \in \mathcal{S}$ is a nilpotent matrix. The *nil index* of \mathcal{S} , $\text{NilInd}(\mathcal{S})$, is defined as the smallest integer k such that $B^k = 0$ for any $B \in \mathcal{S}$. If \mathcal{S} is not nil, then $\text{NilInd}(\mathcal{S}) = \infty$.

A matrix space $\mathcal{S} \leq M(n, \mathbb{F})$ is *nilpotent*, if there exists $k \in \mathbb{N}$, such that for any $B_1, \dots, B_k \in \mathcal{S}$, $\prod_{i \in [k]} B_i = 0$. The smallest such k is called the *nilpotent index* of \mathcal{S} , denoted by $\text{NilptInd}(\mathcal{S})$. If \mathcal{S} is not nilpotent, then $\text{NilptInd}(\mathcal{S}) = \infty$.

Clearly, nilpotent matrix spaces are also nil. Thus $\text{NilInd}(\mathcal{S}) \leq \text{NilptInd}(\mathcal{S})$ for any matrix space \mathcal{S} . The converse is false; for example, it is easy to verify that the space $\left\{ \begin{bmatrix} 0 & x & 0 \\ y & 0 & x \\ 0 & -y & 0 \end{bmatrix} : x, y \in \mathbb{F} \right\}$ is nil but not nilpotent.

The following theorem establishes the basic correspondence between the properties of being acyclic, nilpotent, and nil. In particular, it recovers Theorem 1.8 in Section 1.2.2.

Theorem 4.2. *Let $G = ([n], E)$ be a directed graph, and let $\mathcal{S}_G \leq M(n, \mathbb{F})$ be the graphical matrix space associated with G . Then the following are equivalent:*

- (1) G is acyclic;
- (2) \mathcal{S}_G is nilpotent;
- (3) \mathcal{S}_G is nil.

Furthermore, we have $\text{MaxWalkLen}(G) + 1 = \text{NilInd}(\mathcal{S}_G) = \text{NilptInd}(\mathcal{S}_G)$.

Proof. **(1) \Rightarrow (2):** Suppose G is acyclic. This implies that G has a topological sort, that is, a linear ordering of the vertices such that all edges are directed from left to right. Thus, there exists a permutation matrix P , such that $P^{-1}\mathcal{S}_G P$ is strictly upper-triangular. Then $P^{-1}\mathcal{S}_G P$ is nilpotent, which implies that \mathcal{S}_G is also nilpotent.

(2) \Rightarrow (3): This is straightforward by their definitions.

(3) \Rightarrow (1): We show that if G is cyclic, then \mathcal{S}_G is not nil. Up to relabeling the vertices, suppose the vertex sequence $(1, 2, 3, \dots, k, 1)$ forms a cycle in G , for some $k \geq 1$. Then the following matrix

is in \mathcal{S}_G : $B = \begin{bmatrix} B_0 & 0 \\ 0 & 0 \end{bmatrix}$ where $B_0 \in M(k, \mathbb{F})$ is

$$B_0 = \begin{bmatrix} 0 & 1 & & & \\ & 0 & 1 & & \\ & & \ddots & \ddots & \\ & & & 0 & 1 \\ 1 & & & & 0 \end{bmatrix}_{k \times k}.$$

Clearly, B_0 is not nilpotent. It follows that B is not nilpotent and \mathcal{S}_G is not nil.

The above shows that if G has a cycle, then $\text{MaxWalkLen}(G)$, $\text{NilInd}(\mathcal{S}_G)$ and $\text{NilptInd}(\mathcal{S}_G)$ are all infinite. Thus $\text{MaxWalkLen}(G) + 1 = \text{NilInd}(\mathcal{S}_G) = \text{NilptInd}(\mathcal{S}_G)$ holds whenever G is cyclic.

We now prove that $\text{MaxWalkLen}(G) + 1 = \text{NilptInd}(\mathcal{S}_G) = \text{NilInd}(\mathcal{S}_G)$ when G is acyclic. In this case, $\text{MaxWalkLen}(G)$, $\text{NilInd}(\mathcal{S}_G)$, and $\text{NilptInd}(\mathcal{S}_G)$ are finite.

For $\text{MaxWalkLen}(G) + 1 \geq \text{NilptInd}(\mathcal{S}_G)$. Let $r = \text{MaxWalkLen}(G)$. Note that for a sequence of $r + 1$ elementary matrices, their product is non-zero if and only if their indices form a walk of length $r + 1$. Since there exist no walks of length greater than r in G , any product of a sequence of $r + 1$ such elementary matrices, which come from a linear basis of \mathcal{S}_G , must be zero. It follows that $r + 1 \geq \text{NilptInd}(\mathcal{S}_G)$.

For $\text{NilptInd}(\mathcal{S}_G) \geq \text{NilInd}(\mathcal{S}_G)$. This is straightforward by their definitions.

For $\text{NilInd}(\mathcal{S}_G) \geq \text{MaxWalkLen}(G) + 1$. Let $r = \text{MaxWalkLen}(G)$. We construct a matrix $B \in \mathcal{S}_G$, such that $B^r \neq 0$, which would imply that $\text{NilInd}(\mathcal{S}_G) \geq r + 1$. Since r is finite, every walk in G is a path. By relabeling the vertices, suppose $(1, 2, \dots, r + 1)$ is a path of length r in G . Then the following matrix is in \mathcal{S}_G : $B = \begin{bmatrix} B_0 & 0 \\ 0 & 0 \end{bmatrix}$ where $B_0 \in M(r + 1, \mathbb{F})$ is

$$B_0 = \begin{bmatrix} 0 & 1 & & & & \\ & 0 & 1 & & & \\ & & \ddots & \ddots & & \\ & & & 0 & 1 & \\ & & & & & 0 \end{bmatrix}_{(r+1) \times (r+1)}.$$

Clearly, B_0^r is non-zero, so B^r is non-zero. We then obtain the desired B , concluding the proof. \square

Remark 4.3. Note that G being acyclic implies that the adjacency matrix A_G of G is nilpotent. The reverse direction only holds when A_G is defined over fields of appropriate characteristics. For instance, let G be the n -vertex complete directed graph with self-loops, whose adjacency matrix is the $n \times n$ all-one matrix. Then A_G is nilpotent over any field of characteristic dividing n , although G is cyclic. On the other hand, Theorem 4.2 holds over any field.

4.1.2 Cycle covers and the number of zero eigenvalues

Let G be a directed graph of order n , and $r \in \{0\} \cup [n]$. We say G is r -acyclic if any collection of vertex-disjoint cycles of G covers at most $n - r$ vertices. Note that (1) $r = 0$ corresponds to G having a cycle cover; and (2) $r = n$ corresponds to G being acyclic. The following basic correspondence generalizes Theorem 4.2.

Theorem 4.4. *Let $G = ([n], E)$ be a directed graph and $\mathcal{S}_G \leq M(n, \mathbb{F})$ be the associated graphical matrix space. Then G is r -acyclic if and only if every matrix $B \in \mathcal{S}_G$ has at least r zero eigenvalues.*

Proof. We first show that if every matrix $B \in \mathcal{S}_G$ has at least r zero eigenvalues, then G is r -acyclic. If G is not r -acyclic, we can find disjoint cycles C_1, \dots, C_ℓ which cover $k \geq n - r + 1$ vertices. Let $B \in \mathcal{S}_G$ be the matrix whose (i, j) th entry is 1 if (i, j) is an edge in any of the cycles C_1, \dots, C_ℓ and 0 elsewhere. Since C_1, \dots, C_ℓ are disjoint, B can be written as the direct sum of the adjacency matrices of C_1, \dots, C_ℓ (after a suitable relabeling of the vertices). Then B has $(n - k)$ zero eigenvalues, where $n - k \leq r - 1$, a contradiction.

We then show that if G is r -acyclic, then every $B \in \mathcal{S}_G$ has at least r zero eigenvalues. Note that any matrix $B \in M(n, \mathbb{F})$ having at least r zero eigenvalues is equivalent to that the characteristic

polynomial $P_B(x)$ of B is of the form $x^r Q_B(x)$ for some polynomial $Q_B(x)$ of degree $(n - r)$. On the other hand, write the characteristic polynomial of B as

$$P_B(x) = x^n + (-1)E_1(B)x^{n-1} + \cdots + (-1)^{n-1}E_{n-1}(B)x + (-1)^n E_n(B),$$

where for $k \in [n]$, $E_k(B)$ is the sum of all $k \times k$ principal minors of B (see e.g. [HJ85, Eq. (1.2.13)]). To show B has at least r zero eigenvalues, we need to prove that $E_k(B) = 0$ for all $k \geq n - r + 1$.

Fix $k \geq n - r + 1$, and let $I \subseteq [n]$ be a set of size k . Let B_I be the principal submatrix of B indexed by $I \times I$ and $\mathcal{S}[I] = \{B_I \mid B \in \mathcal{S}\}$. Let $G[I] = (I, E[I])$ be the induced subgraph of G on I , where $E[I] = \{(i, j) \in E \mid i, j \in I\}$. Observe that $\mathcal{S}_G[I] = \mathcal{S}_{G[I]} \leq M(k, \mathbb{F})$. Since G is r -acyclic, then for any I , $G[I]$ cannot be covered by disjoint cycles, implying the bipartite graph corresponding to $G[I]$, i.e., $(I \times I, E[I])$, does not have a perfect matching. By Theorem 1.4, we have every matrix in $\mathcal{S}_{G[I]}$ has rank less than k . This shows that every matrix $B_I \in \mathcal{S}_G[I] = \mathcal{S}_{G[I]}$ has a zero determinant. Further note that

$$E_k(B) = \sum_{I \subseteq [n], |I|=k} \det(B_I).$$

This implies that $E_k(B) = 0$ for any $k \geq n - r + 1$ and any $B \in \mathcal{S}_G$. Thus we conclude that $B \in \mathcal{S}_G$ has at least r zero eigenvalues. \square

Starting from a directed graph $G = ([n], E)$, we can construct a bipartite graph $B(G) = ([n] \times [n], E)$. It is not hard to see that G can be covered by disjoint cycles if and only if $B(G)$ has perfect matchings (which we have used in the second half of the proof). Theorem 4.4 then can be thought of as a directed graph version of Theorem 1.4. The difference (of the statements) is caused by the different underlying symmetries of directed graphs and bipartite graphs: although G and $B(G)$ have the same set of arcs (edges), we may allow different permutations on the left and right vertices in $B(G)$, while there is only one permutation on the vertices of G . Reflecting this fact on the matrix space side, although the directed graph G and the bipartite graph $B(G)$ share the same matrix space \mathcal{S}_G , the symmetries of G and $B(G)$ induce conjugation action and left-right action on \mathcal{S}_G , respectively. Note that the number of zero eigenvalues of $B \in \mathcal{S}_G$, which corresponds to the algebraic multiplicity of the zero eigenvalue of $B \in \mathcal{S}_G$, is invariant under conjugation action; and the rank of $B \in \mathcal{S}_G$, which corresponds to the geometric multiplicity of the zero eigenvalue of $B \in \mathcal{S}_G$, is invariant under left-right action. This justifies the difference between Theorems 1.4 and 4.4.

We remark here again that we expect the basic correspondence in Theorem 4.4 to yield an inherited correspondence, as stated in Conjecture 1.22. However, at the moment, we are unable to prove this. We believe that a proof technique like that of Theorem 1.9 (proved below) may suffice to prove Conjecture 1.22, but there are certain difficulties that we are unable to overcome.

4.2 The inherited correspondence

For a directed graph $G = ([n], E)$, let $\text{MaxAcySize}(G)$ be the maximum size of an acyclic spanning subgraph of G and let $\text{MaxNilDim}(\mathcal{S}_G)$ be the largest dimension of a nil subspace of \mathcal{S}_G . For any subspace \mathcal{S} of $M(n, \mathbb{F})$, its *supporting (directed) graph* $H = ([n], E)$ is defined as follows: $(i, j) \in E$ if there is $A \in \mathcal{S}$ such that $A(i, j) \neq 0$. For two matrix spaces $\mathcal{S}, \mathcal{T} \leq M(n, \mathbb{F})$, $\mathcal{S} \leq \mathcal{T}$ implies the corresponding supporting graph $H_{\mathcal{S}}$ is a spanning subgraph of $H_{\mathcal{T}}$.

The following notion was introduced by de Seguins Pazzis in [dSP13].

Definition 4.5. For a nil matrix space $\mathcal{S} \leq M(n, \mathbb{F})$, we say a non-zero column vector v is \mathcal{S} -adapted, if for any $X \in \mathcal{S}$, $\text{colspan}(X) \neq \langle v \rangle$.

Roughly speaking, the notion of an \mathcal{S} -adapted vector is analogous to the notion of a sink in a directed graph. Indeed, if vertex j in a directed graph G is a sink, then j has no out-neighbors, which means that the j th row of the adjacency matrix of G contains only zeros. Similarly, if the standard column basis vector e_j is \mathcal{S} -adapted, then no non-zero element of \mathcal{S} is supported in the j th row of \mathcal{S} .

An important result in [dSP13] is the following. Both the statement and the proof can be seen as analogs of the basic fact that an acyclic directed graph has a sink.

Lemma 4.6 ([dSP13, Lemma 5]). *For any nil subspace $\mathcal{S} \leq M(n, \mathbb{F})$, there is some $j \in [n]$ such that the j th column standard basis vector e_j is \mathcal{S} -adapted.*

We utilize this lemma to prove the following:

Lemma 4.7. *For any nil subspace $\mathcal{S} \leq M(n, \mathbb{F})$ with supporting graph $H = ([n], F)$, there is an acyclic spanning subgraph $H' = ([n], F')$ of H with $\dim(\mathcal{S}) \leq |F'|$.*

Proof. The lemma holds for $n = 1$. Assume for any nil subspace $\mathcal{S} \leq M(n - 1, \mathbb{F})$ with supporting graph $H = ([n - 1], F)$, such H' exists. Consider a nil subspace $\mathcal{S} \leq M(n, \mathbb{F})$. By Lemma 4.6, without loss of generality, assume e_n is \mathcal{S} -adapted. Let $M \in \mathcal{S}$ be written in the following form

$$M = \begin{bmatrix} A(M) & B(M) \\ C(M) & D(M) \end{bmatrix}$$

where $A(M) \in M(n - 1, \mathbb{F})$, $B(M)^t, C(M) \in \mathbb{F}^{n-1}$ and $D(M) \in \mathbb{F}$. Denote the linear space $A(\mathcal{S}) = \{A(M) : M \in \mathcal{S}\}$. Let $\mathcal{T} = \{M \in \mathcal{S} : B(M) = 0\}$. Since \mathcal{S} is nil, we have $D(M) = 0$ and $A(M)$ is nilpotent for any $M \in \mathcal{T}$. Consider the linear map $\Phi : M(n, \mathbb{F}) \rightarrow M(n, \mathbb{F})$ which maps M to $\begin{bmatrix} 0 & B(M) \\ 0 & 0 \end{bmatrix}$. By the rank-nullity theorem,

$$\dim(\mathcal{S}) = \dim(B(\mathcal{S})) + \dim(\ker(\Phi)|_{\mathcal{S}}) = \dim(B(\mathcal{S})) + \dim(\mathcal{T}).$$

On the other hand, for any matrix $M \in \mathcal{T}$, if $A(M) = 0$, we have $C(M) = D(M) = 0$ since e_n is \mathcal{S} -adapted. This shows that the projection mapping $M \in \mathcal{T}$ to $A(M)$ is injective, which implies that $\dim(\mathcal{T}) = \dim(A(\mathcal{T}))$.

By the induction hypothesis, since $A(\mathcal{T}) \leq M(n - 1, \mathbb{F})$ is nil, there is an acyclic spanning subgraph $H'_{A(\mathcal{T})} = ([n - 1], F'_{A(\mathcal{T})})$ of its supporting graph $H_{A(\mathcal{T})} = ([n - 1], F_{A(\mathcal{T})})$ satisfying $\dim(A(\mathcal{T})) \leq |F'_{A(\mathcal{T})}|$. Note that $H_{A(\mathcal{T})}$ is an induced subgraph of the supporting graph $H = ([n], F)$ of \mathcal{S} . Construct $F_{\mathcal{S}}$ by adding all possible arcs pointing to n to $F'_{A(\mathcal{T})}$ and let $H'_{\mathcal{S}} = ([n], F_{\mathcal{S}})$. Since $H'_{A(\mathcal{T})}$ is acyclic and n is a new vertex, $H'_{\mathcal{S}}$ is also acyclic and $|F_{\mathcal{S}}|$ equals $|F'_{A(\mathcal{T})}|$ plus the in-degree of n in H . Note that the latter is lower bounded by $\dim(B(\mathcal{S}))$. Now we have

$$|F_{\mathcal{S}}| \geq |F'_{A(\mathcal{T})}| + \dim(B(\mathcal{S})) \geq \dim(A(\mathcal{T})) + \dim(B(\mathcal{S})) = \dim(\mathcal{S}). \quad \square$$

Now we are ready to prove Theorem 1.9.

Theorem 1.9, restated. *Let $G = ([n], E)$ be a directed graph, and $\mathcal{S}_G \leq M(n, \mathbb{F})$ be the matrix space associated with G . Then $\text{MaxAcySize}(G) = \text{MaxNilDim}(\mathcal{S}_G)$.*

Proof. It is straightforward to see $\text{MaxAcySize}(G) \leq \text{MaxNilDim}(\mathcal{S}_G)$. Taking a maximum acyclic spanning subgraph $G' = ([n], E')$, we know that $\mathcal{S}_{G'}$ is nil by Theorem 4.2 and $\mathcal{S}_{G'} \leq \mathcal{S}_G$. Thus $\text{MaxAcySize}(G) = \dim(\mathcal{S}_{G'}) \leq \text{MaxNilDim}(\mathcal{S}_G)$.

To see $\text{MaxNilDim}(\mathcal{S}_G) \leq \text{MaxAcySize}(G)$, let \mathcal{S} be the largest nil subspace of \mathcal{S}_G with supporting graph H . By Lemma 4.7, there is an acyclic subgraph H' of H of size at least $\dim(\mathcal{S})$. Noting that H is a spanning subgraph of G , then H' is also a spanning subgraph of G . Thus $\dim(\mathcal{S}) \leq \text{MaxAcySize}(G)$. \square

Remark 4.8. Theorem 1.9 can be thought of as a generalization of Gerstenhaber's theorem on dimensions of nil matrix spaces [Ger58]. Gerstenhaber proved that the largest dimension of a nil matrix spaces is $\frac{n(n-1)}{2}$.⁸ Taking G to be the complete directed graph (with self-loops on every vertex), then $\mathcal{S}_G = \text{M}(n, \mathbb{F})$. In this case, $\text{MaxNilDim}(\text{M}(n, \mathbb{F}))$ is given by the size of the maximum acyclic spanning subgraph of G , which is $\frac{n(n-1)}{2}$. This recovers Gerstenhaber's theorem.

4.3 The induced correspondence

For a directed graph $G = ([n], E)$, let $\text{MaxIndAcyOrd}(G)$ be the maximum order of an acyclic induced subgraph of G . For $\mathcal{S} \leq \text{M}(n, \mathbb{C})$ and $u \leq \mathbb{C}^n$, recall the definition of the induced subspace $\mathcal{S}[U]$ from Definition 1.10. Let $\text{MaxIndNilDim}(\mathcal{S})$ be the largest dimension over U such that $\mathcal{S}[U]$ is nil.

The following lemma is a reformulation of [MOR91, Theorem 4].

Lemma 4.9 ([MOR91]). *Let $\mathcal{S} \leq \text{M}(n, \mathbb{F})$ be a matrix space spanned by rank-1 matrices. Then \mathcal{S} is nilpotent if and only if \mathcal{S} is nil.*

Theorem 1.11, restated. *Let $G = ([n], E)$ be a directed graph, and $\mathcal{S}_G \leq \text{M}(n, \mathbb{C})$ be the associated graphical matrix space. Then $\text{MaxIndAcyOrd}(G) = \text{MaxIndNilDim}(\mathcal{S}_G)$.*

Proof. To see $\text{MaxIndAcyOrd}(G) \leq \text{MaxIndNilDim}(\mathcal{S}_G)$, let $S \subseteq [n]$ be a vertex subset of the maximum order such that the induced subgraph $G[S]$ is acyclic. Let $U = \langle e_i \mid i \in S \rangle$. Then $\mathcal{S}_G[U]$ is nil, as $\mathcal{S}_G[U] = \mathcal{S}_{G[S]}$ and $G[S]$ is acyclic.

To see $\text{MaxIndNilDim}(\mathcal{S}_G) \leq \text{MaxIndAcyOrd}(G)$, let $U \leq \mathbb{C}^n$ be a subspace of the largest dimension such that $\mathcal{S}_G[U]$ is nil. Supposing $\dim(U) = k$, let $T_U \in \text{M}(k \times n, \mathbb{C})$ be a matrix whose rows form an orthonormal basis of U . For $i \in [n]$, let r_i be the i th column of T_U . Then $\mathcal{S}_G[U] = \langle r_i r_j^* \mid (i, j) \in E \rangle$. In particular, $\mathcal{S}_G[U]$ is spanned by rank-1 matrices.

Let $S \subseteq [n]$, $|S| = k$, such that $\{r_i \mid i \in S\}$ is a set of linearly independent column vectors. In particular, for $i \in S$, r_i is non-zero. We claim that $G[S]$ is acyclic. If not, after a possible relabeling of vertices, let $(1, \dots, \ell, 1)$ be a cycle in $G[S]$. Then for any $j \in [\ell - 1]$, $r_j r_{j+1}^*$ is in $\mathcal{S}_G[U]$ and $r_\ell r_1^* \in \mathcal{S}_G[U]$. Note that

$$(r_1 r_2^*)(r_2 r_3^*) \cdots (r_\ell r_1^*) = r_1 (r_2^* r_2) \cdots (r_\ell^* r_\ell) r_1^* = \alpha r_1 r_1^*,$$

where $\alpha = (r_2 r_2^*) \cdots (r_\ell r_\ell^*) \neq 0$. This implies that $(r_1 r_2^*)(r_2 r_3^*) \cdots (r_\ell r_1^*)$ is not nilpotent, so $\mathcal{S}[U]$ is not nilpotent. As $\mathcal{S}[U]$ is spanned by rank-1 matrices, by Lemma 4.9, $\mathcal{S}[U]$ is not nil. This is a contradiction to our assumption, concluding the proof. \square

⁸Gerstenhaber's original proof required the underlying field to be sufficient large [Ger58], this restriction was removed later by [Ser85]. Alternative proofs can be found in, e.g. [MOR91, MMS12, dSP13] and our proof adapts the strategy in [dSP13].

5 Strong connectivity and irreducibility

In this section, we establish the correspondences between strong connectivity and irreducibility. We begin with the basic correspondences, then prove the inherited correspondence in Section 5.2, and finally the induced correspondence in Section 5.3.

5.1 The basic correspondences

Let $G = ([n], E)$ be a directed graph. A set of vertices $V \subseteq [n]$ is *invariant* in G , if there is no arc from V to $[n] \setminus V$. Recall that G is *strongly connected*, if the only non-empty invariant vertex set is $[n]$ itself. A strongly connected component decomposition is a partition $[n] = V_1 \cup \dots \cup V_k$, such that for any $i \in [k]$, $G[V_i]$ is a maximal strongly connected induced subgraph. It is clear that for a graph G , there exists a unique strongly connected component decomposition by choosing $V_1 <_t V_2 <_t \dots <_t V_k$, where $<_t$ is the topological ordering of the strongly connected components of G . Let $c(G)$ be the number of strongly connected components in G .

Let $\mathcal{S} \leq M(n, \mathbb{F})$ be a matrix space. A subspace $U \leq \mathbb{F}^n$ is *invariant* in \mathcal{S} if $\mathcal{S}(U) := \langle \cup_{B \in \mathcal{S}} B(U) \rangle \leq U$. If \mathcal{S} admits a non-trivial invariant subspace (i.e., not $\{0\}$ and \mathbb{F}^n), we say \mathcal{S} is *reducible*; otherwise, we say \mathcal{S} is *irreducible*. Let $0 = U_0 < U_1 < \dots < U_k = \mathbb{F}^n$ be a chain of invariant subspaces of \mathcal{S} . This chain is maximal if for any $i \in [k]$, the induced matrix space of \mathcal{S} on U_i/U_{i-1} is irreducible. By the Jordan–Hölder theorem (see e.g. [Lan05, Theorem 3.5]), two maximal chains of invariant subspaces of \mathcal{S} are of the same length. Let $c(\mathcal{S})$ be the length of a maximal chain of invariant subspaces.

Theorem 5.1. *Let $G = ([n], E)$ be a directed graph, and $\mathcal{S}_G \leq M(n, \mathbb{F})$ be the graphical matrix space associated with G . Then $c(G) = c(\mathcal{S}_G)$. In particular, G is strongly connected if and only if \mathcal{S}_G is irreducible.*

Proof. We begin by proving the special case, that G is strongly connected if and only if \mathcal{S}_G is irreducible (cf. Theorem 1.12). This will then be used in the proof of the more general result that $c(G) = c(\mathcal{S}_G)$.

If G is not strongly connected, then there is a non-empty subset $V \subset [n]$ which is invariant in G . Without loss of generality assume $V = [d]$ for some $d \in [n-1]$. Consider the subspace $U = \langle e_1, \dots, e_d \rangle \leq \mathbb{F}^n$. It is straightforward to verify that U is an invariant subspace. Thus \mathcal{S}_G is not irreducible.

If \mathcal{S}_G is reducible, then there exists a subspace $U \leq \mathbb{F}^n$ of dimension d , $d \in [n-1]$, such that for any $B \in \mathcal{S}_G$, $B(U) \leq U$. Let $T_U \in M(d \times n, \mathbb{F})$ be a matrix whose rows span U . As $\text{rank}(U) = d$, there exist a permutation matrix $P \in GL(n, \mathbb{F})$ and some invertible matrix $T \in GL(d, \mathbb{F})$ such that $\hat{U} = TUP$ is of the form $[I_d \ U_0]$ for some matrix $U_0 \in M(d \times (n-d), \mathbb{F})$. Then the row span of \hat{U} is invariant under the action of $P\mathcal{S}_G P^{-1} = \mathcal{S}_{\hat{G}}$, where \hat{G} is isomorphic to G with respect to the permutation P .

We claim that \hat{G} is not strongly connected, which implies that G is not strongly connected either. To see this, we show that $[d]$ is invariant, namely there is no arc from $[d]$ to $[n] \setminus [d]$. By way of contradiction, suppose there exists $(i, j) \in E(\hat{G})$ with $i \in [d]$ and $j \in [n] \setminus [d]$. Thus $E_{i,j} \in \mathcal{S}_{\hat{G}}$. Note that the i th row vector of \hat{U} is $\hat{u}_i = e_i + (0, u_{0,i}) \in \text{rowspan}(\hat{U})$, where $u_{0,i} \in \mathbb{F}^{n-d}$. The image of \hat{u}_i under $E_{i,j}$ is $\hat{u}_i E_{i,j} = e_j$. As $\text{rowspan}(\hat{U})$ is invariant under $\mathcal{S}_{\hat{G}}$, we have $e_j \in \text{rowspan}(\hat{U})$. This is impossible, because a non-zero vector in $\text{rowspan}(\hat{U})$ has at least one non-zero entry in the first d coordinates. To summarize, there is no arc from $[d]$ to $[n] \setminus [d]$, which shows that \hat{G} is not strongly connected.

To see $c(G) = c(\mathcal{S}_G)$, we first show that $c(G) \leq c(\mathcal{S}_G)$. Suppose $[n] = V_1 \cup \dots \cup V_k$ is the strongly connected component decomposition of G , where $V_1 <_t V_2 <_t \dots <_t V_k$. This naturally gives rise to a chain of subspaces $0 = U_0 < U_1 < U_2 < \dots < U_k = \mathbb{F}^n$ where $U_i = \langle \{e_i \mid i \in V_1 \cup \dots \cup V_i\} \rangle$. It is straightforward to verify that this gives rise to a chain of invariant subspaces of \mathcal{S}_G .

We then show that $c(G) \geq c(\mathcal{S}_G)$. Let $0 = U_0 < U_1 < U_2 < \dots < U_k = \mathbb{F}^n$ be a maximal chain of invariant subspaces of \mathcal{S}_G . Suppose $\dim(U_i) = d_i$. Let $T \in \text{GL}(n, \mathbb{F})$ be a matrix whose first d_i columns span U_i . By a repeated use of Laplace expansion (starting from the last $d_k - d_{k-1}$ columns), there exists $\emptyset = V_0 \subset V_1 \subset V_2 \subset \dots \subset V_k = [n]$, such that for any $i \in [k]$, the submatrix of T with row indices V_i and column indices $[d_i]$ is full-rank. By the argument above, there are no arcs in G that go from V_i to $[n] \setminus V_i$. It follows that for any $i \in [k]$, $V_i \setminus V_{i-1}$ is invariant, and $c(G) \geq c(\mathcal{S}_G)$. \square

5.2 The inherited correspondence

For a directed graph $G = ([n], E)$, let $\text{MaxNscSize}(G)$ be the maximum number of arcs in a non-strongly-connected subgraph of G . Note that $\lambda(G) := |E| - \text{MaxNscSize}(G)$ is known as the *arc-strong connectivity* of G (see e.g. [BJG08, Chap. 1.5]), which is the minimum number of arcs of G whose removal makes G not strongly connected.

For a matrix space $\mathcal{S} \leq \text{M}(n, \mathbb{F})$, let $\text{MaxRdcDim}(\mathcal{S})$ be the largest dimension of a reducible subspace of \mathcal{S} . We are ready to establish the following result:

Theorem 1.13, restated. *Let $G = ([n], E)$ be a directed graph, and $\mathcal{S}_G \leq \text{M}(n, \mathbb{F})$ be the matrix space associated with G . Then $\text{MaxNscSize}(G) = \text{MaxRdcDim}(\mathcal{S}_G)$.*

Proof. To see that $\text{MaxNscSize}(G) \leq \text{MaxRdcDim}(\mathcal{S}_G)$, take a maximum spanning subgraph G' of G which is not strongly connected. Then $\mathcal{S}_{G'} \leq \mathcal{S}_G$ and the reducibility of $\mathcal{S}_{G'}$ follows from Theorem 5.1.

To show that $\text{MaxRdcDim}(\mathcal{S}_G) \leq \text{MaxNscSize}(G)$, suppose that $\mathcal{C} \leq \mathcal{S}_G$ admits a non-trivial and proper invariant subspace $U \leq \mathbb{F}^n$, where $\dim(\mathcal{C}) = \text{MaxRdcDim}(\mathcal{S}_G)$, $\dim(\mathcal{S}_G) = m$, and $\dim(U) = d$. Let $T \in \text{GL}(n, \mathbb{F})$ be an invertible matrix such that $T = \begin{bmatrix} T_U \\ T_V \end{bmatrix}$, where $T_U \in \text{M}(d \times n, \mathbb{F})$ (resp. $T_V \in \text{M}((n-d) \times n, \mathbb{F})$) whose rows span U (resp. a complementary subspace V). Let $T^{-1} = \begin{bmatrix} L & R \end{bmatrix}$ where $L \in \text{M}(n \times d, \mathbb{F})$ and $R \in \text{M}(n \times (n-d), \mathbb{F})$.

Note that for $B \in \text{M}(n, \mathbb{F})$, $TBT^{-1} = \begin{bmatrix} T_UBL & T_UBR \\ T_VBL & T_VBR \end{bmatrix}$. We are interested in $\mathcal{D} := T_U \mathcal{S}_G R \leq \text{M}(d \times (n-d), \mathbb{F})$. Since U is an invariant subspace of $\mathcal{C} \leq \mathcal{S}_G$, for any $C \in \mathcal{C}$, $T_UCR = 0 \in \text{M}(d \times (n-d), \mathbb{F})$. Therefore, for any $C \in \mathcal{C}$, TCT^{-1} is of the form $\begin{bmatrix} C_1 & 0 \\ C_2 & C_3 \end{bmatrix}$ where $C_1 \in \text{M}(d \times d, \mathbb{F})$. It follows from the rank-nullity theorem that $\dim(\mathcal{D}) \leq m - \dim(\mathcal{C}) = m - \text{MaxRdcDim}(\mathcal{S}_G)$. In fact, we have $\dim(\mathcal{D}) = m - \text{MaxRdcDim}(\mathcal{S}_G)$, as otherwise there exists another reducible $\mathcal{C}' \leq \mathcal{S}_G$ of dimension greater than $\text{MaxRdcDim}(\mathcal{S}_G)$ as witnessed by U .

Our goal is to construct a set of arcs in G of size at most $m - \text{MaxRdcDim}(\mathcal{S}_G)$, whose removal makes G not strongly connected. To achieve this, we examine $T = \begin{bmatrix} T_U \\ T_V \end{bmatrix}$. Since $T_U \in \text{M}(d \times n, \mathbb{F})$ is of rank d , there exists a $d \times d$ full-rank submatrix in T_U . Therefore, by a permutation of columns of T and relabeling the vertices of G if necessary, we can put T in the form $\begin{bmatrix} T_{1,1} & T_{1,2} \\ T_{2,1} & T_{2,2} \end{bmatrix}$ where

$T_{1,1} \in \text{GL}(d, \mathbb{F})$. Then T^{-1} is of the form $\begin{bmatrix} S_{1,1} & S_{1,2} \\ S_{2,1} & S_{2,2} \end{bmatrix}$, where $S_{2,2} = T_{2,2} - T_{2,1}T_{1,1}^{-1}T_{1,2}$ is known as

the Schur complement of $T_{1,1}$ [Zha06]. The basic observation that $S_{2,2}$ is invertible will be crucial in the following.

We claim that in $G = ([n], E)$ (after a suitable relabeling), the number of arcs going from $[d]$ to $[n] \setminus [d]$ is no more than $\dim(\mathcal{D}) = m - \text{MaxRdcDim}(\mathcal{S}_G)$. Let $T_{1,1} = [c_1^t \ \dots \ c_d^t]$ where

$c_i \in \mathbb{F}^d$, and $S_{2,2} = \begin{bmatrix} r_1 \\ \vdots \\ r_{n-d} \end{bmatrix}$ where $r_i \in \mathbb{F}^{n-d}$. Since $T_{1,1}$ and $S_{2,2}$ are invertible, c_i 's (resp. r_j 's)

are linearly independent. It follows that $\{T_U E_{i,j} R \mid i \in [d], j \in [n] \setminus [d]\} \subseteq M(d \times (n-d), \mathbb{F})$

are linearly independent, as $T_U E_{i,j} R = [T_{1,1} \ T_{1,2}] E_{i,j} \begin{bmatrix} S_{1,2} \\ S_{2,2} \end{bmatrix} = c_i^t r_{j-d} \in M(d \times (n-d), \mathbb{F})$. Let

the arcs in $G = ([n], E)$ from $[d]$ to $[n] \setminus [d]$ be $(i_1, j_1), \dots, (i_\ell, j_\ell)$, where $i_1, \dots, i_\ell \in [d]$ and $j_1, \dots, j_\ell \in [n] \setminus [d]$. Then \mathcal{D} contains $T_U E_{i_k, j_k} R$, $k \in [\ell]$, which are linearly independent. It follows that $\ell \leq \dim(\mathcal{D}) = m - \text{MaxRdcDim}(\mathcal{S}_G)$. That is, by removing $\ell \leq m - \text{MaxRdcDim}(\mathcal{S}_G)$ arcs, G becomes not strongly connected. This concludes the proof. \square

5.3 The induced correspondence

For a directed graph $G = ([n], E)$, let $\text{MaxIndNscOrd}(G)$ be the maximum order of a non-strongly-connected induced subgraph of G . Note that $\kappa(G) := n - \text{MaxIndNscOrd}(G)$ is known as the *vertex-strong connectivity* of G (see e.g. [BJG08, Chap. 1.5]), which is the minimum number of vertices of G whose removal makes G not strongly connected.

For a matrix space $\mathcal{S} \leq M(n, \mathbb{C})$, let $\text{MaxIndRdcDim}(\mathcal{S})$ be the largest dimension of $U \leq \mathbb{C}^n$ such that $\mathcal{S}[U]$ is reducible. We are ready to establish the following result:

Theorem 1.14, restated. *Let $G = ([n], E)$ be a directed graph, and $\mathcal{S}_G \leq M(n, \mathbb{C})$ be the matrix space associated with G . Then $\text{MaxIndNscOrd}(G) = \text{MaxIndRdcDim}(\mathcal{S}_G)$.*

Proof. To see that $\text{MaxIndNscOrd}(G) \leq \text{MaxIndRdcDim}(\mathcal{S}_G)$, take a maximum induced subgraph $G' = (V(G'), E(G'))$ of G which is not strongly connected. Let U be the subspace spanned by $\{e_i \mid i \in V(G')\}$. Then by Theorem 5.1, $\mathcal{S}_G[U] \leq M(|V(G')|, \mathbb{C})$ is reducible since G' is not strongly connected.

To see that $\text{MaxIndNscOrd}(\mathcal{S}_G) \leq \text{MaxIndRdcDim}(G)$, take a subspace U of dimension d such that $\mathcal{S}_G[U]$ admits a non-trivial invariant subspace $V \leq U$. Let $T_U \in M(d \times n, \mathbb{C})$ be the matrix whose rows are an orthonormal basis of U . Take a unitary matrix $R \in M(d, \mathbb{C})$ which maps the first $b = \dim(V)$ orthonormal basis vectors of U to an orthonormal basis of V . Then for any matrix $B \in \mathcal{S}_G$,

$$RT_U B T_U^* R^* = \begin{bmatrix} B_{1,1} & 0 \\ B_{2,1} & B_{2,2} \end{bmatrix} \quad (5.1)$$

for some $B_{1,1} \in M(b, \mathbb{C})$, $B_{2,1} \in M(b \times (d-b), \mathbb{C})$ and $B_{2,2} \in M(d-b, \mathbb{C})$. Without loss of generality, assume the first d columns of $RT_U \in M(d \times n, \mathbb{C})$ are linearly independent and denote this submatrix as $T \in GL(d, \mathbb{C})$. (Otherwise, we can permute the columns of T_U and permute the vertices of G accordingly). Moreover, there is a partition $[d] = I_1 \cup I_2$ with $|I_1| = b$, such that the submatrix T_1 (resp. T_2) of T with column indices from I_1 (resp. I_2) and row indices from $[b]$ (resp. $[n] \setminus [b]$) is invertible.

We shall prove that there is no edge of $G[[d]]$ going from I_1 to I_2 . Suppose for contradiction that we have $(i, j) \in E$ for some $i \in I_1$ and $j \in I_2$. Let the i th column of RT_U be t_i . Then $RT_U E_{i,j} T_U^* R^* = t_i t_j^*$ is of the form as in Eq. (5.1). Thus, the k th coordinate of t_i is 0 for all $k \in [b]$ or the ℓ th coordinate of t_j is 0 for all $\ell \in [d] \setminus [b]$. In the first case, we have the i th column of T_1

is 0, contradicting T_1 being invertible; and in the second case, we have the j th column of T_2 is 0, contradicting T_2 being invertible. Thus there cannot be an edge from I_1 to I_2 , which proves that $G[[d]]$ is not strongly connected, and $\text{MaxIndNscOrd}(\mathcal{S}_G) \leq \text{MaxIndRdcDim}(G)$ follows. \square

6 Isomorphism, congruence, and conjugacy

Recall that two directed graphs $G = ([n], E)$ and $H = ([n], F)$ are isomorphic, if there exists a bijective map $f : [n] \rightarrow [n]$, such that $(i, j) \in E$ if and only if $(f(i), f(j)) \in F$. In this section we extend this notion to graphical matrix spaces. Recall that two matrix spaces $\mathcal{S}, \mathcal{C} \leq \text{M}(n, \mathbb{F})$ are conjugate (resp. congruent) if and only if there exists $T \in \text{GL}(n, \mathbb{F})$, such that $\mathcal{S} = TCT^{-1}$ (resp. $\mathcal{S} = TCT^t$). When $\mathbb{F} = \mathbb{C}$, $\mathcal{S}, \mathcal{C} \leq \text{M}(n, \mathbb{F})$ are congruent if and only if there exists $T \in \text{GL}(n, \mathbb{F})$, such that $\mathcal{S} = TCT^*$.

Note that the action of S_n , when considered as a subgroup of $\text{GL}(n, \mathbb{F})$, is compatible with both the conjugacy and congruence actions of $\text{GL}(n, \mathbb{F})$. This immediately gives us the following.

Observation 6.1. *Let $G = ([n], E)$ and $H = ([n], F)$ be two directed graphs and \mathcal{S}_G and \mathcal{S}_H be the associated graphical matrix spaces, respectively. If G is isomorphic to a subgraph of H , then \mathcal{S}_G is conjugate to (resp. congruent to) a subspace of \mathcal{S}_H .*

6.1 Isomorphism and congruence

We show that graphical matrix space congruence implies graph isomorphism by proving a stronger result: G is isomorphic to a subgraph of H if and only if \mathcal{S}_G is congruent to a subspace of \mathcal{S}_H .

Proposition 6.2. *Let $G = ([n], E)$ and $H = ([n], F)$ be directed graphs. Let \mathcal{S}_G and \mathcal{S}_H be the associated graphical matrix spaces, respectively. If \mathcal{S}_G is congruent to a subspace of \mathcal{S}_H , then G is isomorphic to a subgraph of H .*

Proof. Let $T \in \text{GL}(n, \mathbb{F})$ such that $T\mathcal{S}_G T^t \leq \mathcal{S}_H$. Let $t_{i,j} \in \mathbb{F}$ be the (i, j) th entry of T . Then for any $(i, j) \in E$, we have

$$TE_{i,j}T^t = \begin{bmatrix} t_{1,i} \\ \vdots \\ t_{n,i} \end{bmatrix} [t_{1,j} \quad \cdots \quad t_{n,j}] = [t_{k,i}t_{\ell,j}]_{k,\ell \in [n]} \in \mathcal{S}_H.$$

By the structure of \mathcal{S}_H , if the (k, ℓ) th entry $t_{k,i}t_{\ell,j}$ of $TE_{i,j}T^t$ is non-zero, then $(k, \ell) \in F$. As T is invertible, there exists a permutation $\sigma : [n] \rightarrow [n]$ such that $t_{\sigma(i),i} \neq 0$ for any $i \in [n]$. In particular, $t_{\sigma(i),i}t_{\sigma(j),j} \neq 0$ for any $i, j \in [n]$. Thus, we have that for $(i, j) \in E$, $(\sigma(i), \sigma(j)) \in F$. In other words, σ is an injective map from vertices of G to vertices of H which preserves arcs, i.e., an embedding of G into H . This shows that G is a subgraph of H , as claimed. \square

Remark 6.3. Recall that when $\mathbb{F} = \mathbb{C}$, we define congruence using T^* rather than T^t . It can be verified that the above argument works verbatim in case $\mathbb{F} = \mathbb{C}$, by simply replacing every instance of a transpose by a conjugate transpose and by replacing $t_{\ell,j}$ by $\overline{t_{\ell,j}}$.

The proof of Proposition 6.2 yields the following corollary which will be useful in Section 7.

Corollary 6.4. *Let $G = ([n], E)$ and $H = ([n], F)$ be directed graphs with associated graphical matrix spaces \mathcal{S}_G and \mathcal{S}_H , respectively. If $T\mathcal{S}_G T^t = \mathcal{S}_H$ for some $T \in \text{GL}(n, \mathbb{F})$, then any permutation σ satisfying that the $(i, \sigma(i))$ th entry of T is non-zero for each $i \in [n]$ gives an isomorphism between G and H .*

The next natural thing to examine is whether this basic correspondence can be boosted to an inherited correspondence. Namely, is the maximum size over subgraphs of G that do not contain an isomorphic copy of H equals the largest dimension over subspaces of \mathcal{S}_G that do not contain a congruent copy of \mathcal{S}_H ? Interestingly, this is not true in general, as seen in the following example.

Example 6.5. Let $G = ([2], \{(1, 2), (2, 1)\})$ and $H = ([2], \{(1, 2)\})$. Then any non-empty subgraph of G contains an isomorphic copy of H , thus the maximum size over subgraphs of G with no isomorphic copy of H is 0. On the other hand, take the one-dimensional subspace $\mathcal{S} = \left\{ \begin{bmatrix} 0 & x \\ x & 0 \end{bmatrix} : x \in \mathbb{F} \right\}$ of \mathcal{S}_G . As congruence preserves symmetry, \mathcal{S}_H cannot be a subspace of any congruent copy of \mathcal{S} . Then the largest dimension over subspaces of \mathcal{S}_G that do not contain a congruent copy of \mathcal{S}_H is at least 1.

6.2 Isomorphism and conjugacy

We now show that graphical matrix space conjugacy implies graph isomorphism. The proof is inspired by [BS20, Theorem 4.13 in arXiv version 2] and is much more complicated than that of Proposition 6.2. To prepare for its proof, we need the following notation.

A 3-way array is a rectangular cuboid of field elements. Let $\mathbb{T}(\ell \times n \times m, \mathbb{F})$ be the linear space of $\ell \times n \times m$ 3-way arrays over \mathbb{F} . We use the fixed-width teletype font for 3-way arrays, such as \mathbf{A} , \mathbf{B} , and so on.

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

A 3-way array allows for general linear group actions in three directions. Given $P \in \mathbb{M}(\ell, \mathbb{F})$ and $Q \in \mathbb{M}(n, \mathbb{F})$, let \mathbf{PAQ} be the $\ell \times n \times m$ 3-way array whose k th frontal slice is $\mathbf{PA}_k\mathbf{Q}$. For $R = (r_{i,j}) \in \text{GL}(m, \mathbb{F})$, let \mathbf{A}^R be the $\ell \times n \times m$ 3-way array whose k th frontal slice is $\sum_{k' \in [m]} r_{k',k} A_{k'}$.

Given $\mathbf{A} \in \mathbb{T}(\ell \times n \times m, \mathbb{F})$, sometimes it is convenient to work with the associated 3-way arrays obtained by permuting the three indices. For example, from \mathbf{A} above we can construct a 3-way array $\mathbf{B} \in \mathbb{T}(n \times m \times \ell, \mathbb{F})$, such that the i th frontal slice of \mathbf{B} is the i th horizontal slice of \mathbf{A} .

Let $\mathcal{A}, \mathcal{B} \leq \mathbb{M}(n, \mathbb{F})$ and suppose $\dim(\mathcal{A}) = \dim(\mathcal{B}) = m$. We can form 3-way arrays \mathbf{A} and \mathbf{B} in $\mathbb{T}(n \times n \times m, \mathbb{F})$, by taking ordered bases of \mathcal{A} and \mathcal{B} , respectively. Then \mathbf{A} and \mathbf{B} are conjugate, if and only if there exist $T \in \text{GL}(n, \mathbb{F})$ and $R \in \text{GL}(m, \mathbb{F})$, such that $\mathbf{TAT}^{-1} = \mathbf{B}^R$.

We now turn to graphical matrix spaces, and deduce a key lemma (Lemma 6.8) that characterizes conjugacy of graphical matrix spaces. The following notation will be used several times in slightly different contexts, so we collect them as a definition for future references.

Definition 6.6. Let $G = ([n], E)$ and $H = ([n], F)$ be directed graphs of size m . Let \mathcal{S}_G and \mathcal{S}_H be the associated graphical matrix spaces, respectively. Construct $\mathbf{S}_G \in \mathbb{T}(n \times n \times m, \mathbb{F})$ by setting its frontal slices as $(E_{i,j} \mid (i, j) \in E)$, where edges in G are ordered lexicographically. Similarly construct $\mathbf{S}_H \in \mathbb{T}(n \times n \times m, \mathbb{F})$. Let $\mathbf{C} \in \mathbb{T}(n \times m \times n, \mathbb{F})$ be the 3-way array whose frontal slices are the horizontal slices of \mathbf{S}_G . That is, $\mathbf{C}(i, j, k) = \mathbf{S}_G(k, i, j)$. Similarly, let $\mathbf{D} \in \mathbb{T}(n \times m \times n, \mathbb{F})$ be the 3-way array whose frontal slices are the horizontal slices of \mathbf{S}_H . Let C_i 's be the frontal slices of \mathbf{C} , and D_i 's the frontal slices of \mathbf{D} .

Observation 6.7. Let C_i and D_i be as in Definition 6.6. The non-zero columns of C_i 's are non-overlapping. That is, for any $k \in [m]$, there exists one and only one $i \in [n]$, such that the k th column

of C_i is non-zero. The above statement holds for D_i 's too. In particular, for any $T \in \text{GL}(n, \mathbb{F})$, the non-zero columns of TD_i 's are also non-overlapping.

Proof. Suppose the k th edge by lexicographic order in G is (i, j) . So the k th frontal slice of \mathcal{S}_G is $E_{i,j}$. The k th column of C_i is the transpose of the i th row of $E_{i,j}$, namely e_j . For any $i' \neq i$, the k th column of $C_{i'}$ is the transpose of the i' th row of $E_{i,j}$, which is the all-zero vector. The same argument works for D_i , and therefore TD_i too. \square

An important consequence of Observation 6.7 is the following.

Lemma 6.8. *Let $\mathcal{S}_G, \mathcal{S}_H \leq \text{M}(n, \mathbb{F})$ with $\dim(\mathcal{S}_G) = \dim(\mathcal{S}_H) = m$, $\mathbf{S}_G, \mathbf{S}_H \in \text{T}(n \times n \times m, \mathbb{F})$, and $\mathbf{C}, \mathbf{D} \in \text{T}(n \times m \times n, \mathbb{F})$ be in Definition 6.6. Then \mathcal{S}_G and \mathcal{S}_H are conjugate, if and only if there exists $T \in \text{GL}(n, \mathbb{F})$, such that the following holds: letting C'_i be the i th frontal slice of \mathbf{C}^T , then for any $i \in [n]$, $\text{colspan}(C'_i) \subseteq \text{colspan}(TD_i)$.*

Proof. For the only if direction, note that \mathcal{S}_G and \mathcal{S}_H being conjugate implies that there exist $T \in \text{GL}(n, \mathbb{F})$ and $R \in \text{GL}(m, \mathbb{F})$, such that $T^t \mathbf{S}_G = \mathbf{S}_H^R T^t$. (The use of T^t is to make some notation in the following easier.) This translates to $\mathbf{C}^T = TDR$. Therefore, for any $i \in [n]$, $C'_i = TD_i R$, which implies that $\text{colspan}(C'_i) \subseteq \text{colspan}(TD_i)$.

For the if direction, by Observation 6.7, the non-zero columns of TD_i over $i \in [n]$ are non-overlapping. Let $[m] = S_1 \cup \dots \cup S_n$ be a partition such that the indices of non-zero columns in TD_i are exactly in S_i . Suppose $TD_i = [u_1, \dots, u_m]$, where $u_i \in \mathbb{F}^n$. Let $R = (r_{i,j})_{i,j \in [m]} \in \text{M}(m, \mathbb{F})$. Then the k th column of $TD_i R$ is $\sum_{j \in [m]} r_{j,k} u_j = \sum_{j \in S_i} r_{j,k} u_j$, where the equality is due to the fact that $u_j = 0$ for $j \notin S_i$. In other words, for any $i \in [n]$, only the columns of R whose indices in S_i are effective for the columns TD_i . Using the condition that $\text{colspan}(C'_i) \subseteq \text{colspan}(TD_i)$, we can set the rows of $R \in \text{M}(m, \mathbb{F})$ with indices in S_i , such that $C'_i = TD_i R$, for each $i \in [n]$. That is, $\mathbf{C}^T = TDR$, which implies that $T^t \mathbf{S}_G = \mathbf{S}_H^R T^t$. Now we recall that the frontal slices of \mathcal{S}_G and \mathcal{S}_H are bases of \mathcal{S}_G and \mathcal{S}_H which are both m -dimensional. It follows that R must be invertible, and $T^t \mathbf{S}_G T^{-t} = \mathbf{S}_H$. \square

We are now ready to show that conjugacy implies isomorphism.

Proposition 6.9. *Let $G = ([n], E)$ and $H = ([n], F)$ be directed graphs of size m . Let \mathcal{S}_G and \mathcal{S}_H be the associated graphical matrix spaces, respectively. If \mathcal{S}_G and \mathcal{S}_H are conjugate, then G and H are isomorphic.*

Proof. From \mathcal{S}_G and \mathcal{S}_H , we construct $\mathbf{S}_G, \mathbf{S}_H \in \text{T}(n \times n \times m, \mathbb{F})$ and $\mathbf{C}, \mathbf{D} \in \text{T}(n \times m \times n, \mathbb{F})$ as in Definition 6.6. Since \mathcal{S}_G and \mathcal{S}_H are conjugate, there exist $T \in \text{GL}(n, \mathbb{F})$ and $R \in \text{GL}(m, \mathbb{F})$ such that $T^t \mathbf{S}_G = \mathbf{S}_H^R T^t$, which yields $\mathbf{C}^T = TDR$.

Denote $T = [t_1 \ t_2 \ \dots \ t_n]$, where $t_i = (t_{i,1}, \dots, t_{i,n})^t$ is the i th column. By the Laplace expansion, there exists $i \in [n]$, such that $t_{1,i} \neq 0$ and $T' = [e_i \ t_2 \ \dots \ t_n]$ is also invertible.

Claim 6.10. *Let C'_i be the i th frontal slice of $\mathbf{C}^{T'}$. Then for any $i \in [n]$, $\text{colspan}(C'_i) \subseteq \text{colspan}(T' D_i)$.*

Once Claim 6.10 holds, it would follow that $T'^t \mathbf{S}_G T'^{-t} = \mathbf{S}_H$ by Lemma 6.8. We could then apply this procedure to T' to set t_2 to be a standard basis vector, and so on, until we get a permutation matrix P such that $P^{-1} \mathbf{S}_G P = \mathbf{S}_H$. This would allow us to conclude that G and H are isomorphic.

We now prove Claim 6.10.

Proof of Claim 6.10. As the difference between T and T' lies on the first column, changing from T to T' has the following possible consequences:

1. For \mathfrak{C}^T and $\mathfrak{C}^{T'}$, their first frontal slices are different.
2. For those frontal slices D_j containing e_1 as a column, the j th frontal slices of TD and $T'D$ are different.

We now distinguish the following cases.

- (a) Suppose $e_1 \notin \text{colset}(D_j)$ for any $j \geq 2$. This is when neither change happens. Consider $\sum_{k \in [n]} t_{k,j} C_k$, which is equal to $TD_j R$. It is clear that $\text{colspan}(C'_j) = \text{colspan}(\sum_{k \in [n]} t_{k,j} C_k) = \text{colspan}(TD_j) = \text{colspan}(T'D_j)$.
- (b) Suppose $e_1 \in \text{colset}(D_j)$ for some $j \geq 2$. Then $t_1 \in \text{colset}(TD_j)$. Consider $\sum_{k \in [n]} t_{k,j} C_k$, which is equal to $TD_j R$. That is, $\text{colspan}(\sum_{k \in [n]} t_{k,j} C_k) = \text{colspan}(TD_j)$. By Observation 6.7, the columns of $\sum_{k \in [n]} t_{k,j} C_k$ are scaled standard basis vectors. As the i th entry of $t_1 \in \text{colspan}(\sum_{k \in [n]} t_{k,j} C_k)$ is non-zero, it follows that the i th standard basis vector e_i is in $\text{colset}(\sum_{k \in [n]} t_{k,j} C_k)$. It follows that $\text{colspan}(\sum_{k \in [n]} t_{k,j} C_k) \supseteq \text{colspan}(T'D_j)$. Since T and T' are both full-rank, we have $\dim(\text{colspan}(\sum_{k \in [n]} t_{k,j} C_k)) = \dim(\text{colspan}(TD_j)) = \dim(\text{colspan}(T'D_j))$. We then have $\text{colspan}(C'_j) = \text{colspan}(\sum_{k \in [n]} t_{k,j} C_k) = \text{colspan}(T'D_j)$, where the first equality is due to the assumption that $j \geq 2$.
- (c) Suppose $e_1 \notin \text{colset}(D_1)$. The first frontal slice of \mathfrak{C}^T is $\sum_{k \in [n]} t_{k,1} C_k = TD_1 R$. Note that $t_{i,1} \neq 0$, and $C'_1 = C_i$. We have $\text{colspan}(C'_1) = \text{colspan}(C_i) \subseteq \text{colspan}(\sum_{k \in [n]} t_{k,1} C_k) = \text{colspan}(TD_1) = \text{colspan}(T'D_1)$, where the inclusion is due to Lemma 6.8 and the last equality is due to the assumption that $e_1 \notin \text{colset}(D_1)$.
- (d) Suppose $e_1 \in \text{colset}(D_1)$. This is when both changes happen: the first frontal slice changes from $\sum_{k \in [n]} t_{k,1} C_k$ to C_i , and $t_1 \in \text{colset}(TD_1)$ changes to $e_i \in \text{colset}(T'D_1)$. Still, we have $\text{colspan}(\sum_{k \in [n]} t_{k,1} C_k) = \text{colspan}(T'D_1)$ by the argument for case (b), and $\text{colspan}(C'_1) = \text{colspan}(C_i) \subseteq \text{colspan}(\sum_{k \in [n]} t_{k,1} C_k)$ by the argument for case (c). These allow us to deduce that $\text{colspan}(C'_1) \subseteq \text{colspan}(T'D_1)$. \square

The proof of Claim 6.10 concludes the proof of Proposition 6.9. \square

The proof of Proposition 6.9 also yields the following.

Corollary 6.11. *Let $G = ([n], E)$ and $H = ([n], F)$ be directed graphs. Let \mathcal{S}_G and \mathcal{S}_H be the corresponding graphical matrix spaces, respectively. Let $T = (t_{i,j}) \in \text{GL}(n, \mathbb{F})$ satisfy that $T\mathcal{S}_G T^{-1} = \mathcal{S}_H$. For $i, j \in [n]$, let $T_{(i,j)}$ be the submatrix of T obtained by deleting the i th row and the j th column. If $t_{i,j} \neq 0$ and $T_{(i,j)}$ is invertible, then there exists an isomorphism π from G to H , such that $\pi(i) = j$.*

In contrast to Proposition 6.2, there is no embedding version of this correspondence: the following example illustrates that there exist directed graphs G and H such that G is not isomorphic to any subgraph of H while \mathcal{S}_G is conjugate to a subspace of \mathcal{S}_H .

Example 6.12. Let $G = ([3], \{(1, 1), (2, 2)\})$ and $H = ([3], [3] \times [3] \setminus \{(1, 1), (3, 3)\})$. Their graphical matrix spaces are

$$\mathcal{S}_G = \left\{ \begin{bmatrix} a & 0 & 0 \\ 0 & b & 0 \\ 0 & 0 & 0 \end{bmatrix} : a, b \in \mathbb{F} \right\} \text{ and } \mathcal{S}_H = \left\{ \begin{bmatrix} 0 & x_1 & x_2 \\ y_1 & z & x_3 \\ y_2 & y_3 & 0 \end{bmatrix} : x_1, x_2, x_3, y_1, y_2, y_3, z \in \mathbb{F} \right\}.$$

As G has two self-loops while H only has one, G is not isomorphic to any subgraph of H . However, setting $T = \begin{bmatrix} 1 & 0 & 1 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{bmatrix} \in \text{GL}(3, \mathbb{F})$, we have

$$T\mathcal{S}_G T^{-1} = \left\{ \begin{bmatrix} 0 & a & -a \\ -b & a+b & -a \\ -b & b & 0 \end{bmatrix} : a, b \in \mathbb{F} \right\},$$

which is a subspace of \mathcal{S}_H .

7 Vertex transitivity and conjugacy/congruence irreducibility

Let $G = ([n], E)$ be a directed graph. Let $\text{Aut}(G) \leq S_n$ be the automorphism group of G . Recall that G is *vertex-transitive*, if $\text{Aut}(G)$ is a transitive group. Let $\mathcal{S} \leq M(n, \mathbb{F})$. Let $\text{Conj}(\mathcal{S}) := \{T \in \text{GL}(n, \mathbb{F}) \mid T\mathcal{S}T^{-1} = \mathcal{S}\} \leq \text{GL}(n, \mathbb{F})$. We say that \mathcal{S} is *conjugacy irreducible*, if $\text{Conj}(\mathcal{S})$ is irreducible as a matrix group. Let $\text{Cong}(\mathcal{S}) := \{T \in \text{GL}(n, \mathbb{F}) \mid T\mathcal{S}T^t = \mathcal{S}\} \leq \text{GL}(n, \mathbb{F})$. We say that \mathcal{S} is *congruence irreducible*, if $\text{Cong}(\mathcal{S})$ is irreducible as a matrix group.

Proposition 7.1. *Let \mathbb{F} be a field of order > 2 . Let $G = ([n], E)$ be a directed graph, and $\mathcal{S}_G \leq M(n, \mathbb{F})$ be the matrix space associated with G . If G is vertex-transitive, then \mathcal{S}_G is conjugacy irreducible and congruence irreducible.*

Proof. Let $P = \text{Aut}(G) \leq S_n$ be a transitive group. Embed P as a subgroup of $\text{GL}(n, \mathbb{F})$, and let D be the group of invertible diagonal matrices in $\text{GL}(n, \mathbb{F})$. By $P, D \leq \text{Conj}(\mathcal{S}_G)$ (resp. $P, D \leq \text{Cong}(\mathcal{S}_G)$), the group $H = \langle P, D \rangle$ is in $\text{Conj}(\mathcal{S}_G)$ (resp. $\text{Cong}(\mathcal{S}_G)$). Note that H is a subgroup of the monomial group. We claim that H is irreducible. To see this, by the action of D and $|\mathbb{F}| > 2$,⁹ the only possible invariant subspaces are sums of coordinate subspaces. Then the action of P ensures that these invariant subspaces cannot be proper. This proves the claim and implies \mathcal{S}_G is conjugacy irreducible and congruence irreducible. \square

Now we prove that conjugacy and congruence irreducibility imply vertex-transitivity. To this end, we first show the graphs are regular and relate the automorphisms of graphs to the invertible matrices in $\text{Conj}(\mathcal{S}_G)$ and $\text{Cong}(\mathcal{S}_G)$. First of all, we recall some notation. Let $\mathbf{S}_G \in T(n \times n \times m, \mathbb{F})$, $\mathbf{C} \in T(n \times m \times n, \mathbb{F})$, and $C_i \in M(n \times m, \mathbb{F})$ be as in Definition 6.6. Note that C_i records the information of the out-neighbors of the vertex i , and $\text{rank}(C_i) = \dim(\text{colspan}(C_i))$ is equal to

the out-degree of i . Let $T \in \text{Conj}(\mathcal{S}_G)$ (resp. $T \in \text{Cong}(\mathcal{S}_G)$) and write $T = \begin{bmatrix} t_1 \\ \vdots \\ t_n \end{bmatrix}$, where $t_i =$

$(t_{i,1}, \dots, t_{i,n})$ is the i th row of T . Then $\mathbf{C}^{T^t} = T^t \mathbf{C} R$ (resp. $\mathbf{C}^{T^t} = T^{-1} \mathbf{C} R$) for some $R \in \text{GL}(m, \mathbb{F})$. Denote by C'_i the i th frontal slice of \mathbf{C}^{T^t} .

Lemma 7.2. *Let $G = ([n], E)$ be a directed graph, and $\mathcal{S}_G \leq M(n, \mathbb{F})$ be the matrix space associated with G . If \mathcal{S}_G is conjugacy irreducible or congruence irreducible, then G is regular.*

Proof. We first show that \mathcal{S}_G being conjugacy irreducible (resp. congruence irreducible) implies that G is out-regular, i.e., every vertex is of the same out-degree. By way of contradiction, let $S \subseteq [n]$ be the set of vertices of the smallest out-degree d . Then for any given $i \in S$, $\dim(\text{colspan}(T^t C_i)) =$

⁹If $|\mathbb{F}| = 2$, H is reducible because of the all-one vector.

$\text{rank}(C_i) = d$ (resp. $\dim(\text{colspan}(T^{-1}C_i)) = \text{rank}(C_i) = d$). By Lemma 6.8 (resp. Lemma 6.8 with changing the conclusion to $\text{colspan}(C'_i) \subseteq \text{colspan}(T^{-t}D_i)$), $\dim(\text{colspan}(C'_i)) \leq d$.

We claim that the indices of the non-zero entries of t_i must be in S , i.e., $t_{i,j} = 0$ for any $j \notin S$. To see this, recall that $C'_i = \sum_{k=1}^n t_{i,k}C_k$. So if for some $j \notin S$, $t_{i,j} \neq 0$, then this C_j with $\text{rank}(C_j) > d$ would be involved in C'_i , which would result in $\dim(\text{colspan}(C'_i)) > d$ by Observation 6.7.

Thus, we have $t_{i,j} = 0$ for any $i \in S$ and $j \notin S$. It follows that $\langle e_i, i \in S \rangle$ is an invariant subspace of T , giving the desired contradiction.

Recall that C_i 's are the horizontal slices of \mathcal{S}_G . Applying the same argument to the vertical slices of \mathcal{S}_G , we can deduce that G is also in-regular, i.e., every vertex is of the same in-degree. It follows G is regular. \square

Lemma 7.3. *Let $G = ([n], E)$ be a directed graph and $\mathcal{S}_G \leq M(n, \mathbb{F})$ be the matrix space associated with G . If \mathcal{S}_G is conjugacy irreducible (resp. congruence irreducible), then for $T \in \text{Conj}(\mathcal{S}_G)$ (resp. $T \in \text{Cong}(\mathcal{S}_G)$) and $i, j \in [n]$, we have*

- (1) *vertices k and ℓ share the same out-neighborhood if $t_{i,k}$ and $t_{i,\ell}$ are non-zero;*
- (2) *vertices k and ℓ share the same in-neighborhood if $t_{k,j}$ and $t_{\ell,j}$ are non-zero.*

Proof. We first prove (1) by contradiction. By Lemma 7.2, G is d -out-regular, which implies that $\dim(\text{colspan}(T^t C_i)) = \text{rank}(C_i) = d$ (resp. $\dim(\text{colspan}(T^{-1}C_i)) = \text{rank}(C_i) = d$) for any $i \in [n]$. Now assume that vertex k and ℓ do not share the same out-neighborhood. Thus, $\text{colspan}(C_k) \neq \text{colspan}(C_\ell)$. Note that $C'_i = \sum_{j=1}^n t_{i,j}C_j$ where $t_{i,k}$ and $t_{i,\ell}$ are non-zero. (Recall that C'_i the i th frontal slice of \mathcal{C}^{T^t} .) By Observation 6.7, we have $\dim(\text{colspan}(C'_i)) > d$. This leads to a contradiction to the statement of Lemma 6.8 (resp. Lemma 6.8 with changing the conclusion to $\text{colspan}(C'_i) \subseteq \text{colspan}(T^{-t}D_i)$). The proof for (2) is the same by considering the vertical slices. \square

We are now ready to prove that conjugacy irreducibility or congruence irreducibility implies vertex-transitivity.

Proposition 7.4. *Let \mathbb{F} be a field of order > 2 . Let $G = ([n], E)$ be a directed graph, and $\mathcal{S}_G \leq M(n, \mathbb{F})$ be the matrix space associated with G . If \mathcal{S}_G is conjugacy irreducible, G is vertex transitive.*

Proof. Recall that $T = (t_{i,j}) \in \text{Conj}(\mathcal{S}_G)$. We use $T_{(i,j)}$ to denote the matrix by deleting the i th row and j th column from T , and $T_{(i,j),(i',j')}$ to denote the matrix by deleting the i th and i' th rows as well as j th and j' th columns from T , and so on.

By way of contradiction, suppose G is not vertex-transitive. Then the action of $\text{Aut}(G) \leq S_n$ on $[n]$ has at least two orbits. Let $S \subset [n]$ be one orbit of $\text{Aut}(G)$. By a permutation of the vertices if necessary, we can set $S = [s]$ for $s = |S|$.

Claim 7.5. *For any $i \in [s]$ and $j \in [n] \setminus [s]$, $t_{i,j} = 0$.*

Proof. By contradiction and without loss of generality, assume $t_{1,j} \neq 0$ for some $j \geq s + 1$. Note that no automorphism of G can send any $i \in [s]$ to any $j \in [n] \setminus [s]$ and vice versa. By Corollary 6.11, $T_{(1,j)}$ must be singular. On the other hand, since T is invertible, there exists $i \in [n]$ such that $t_{1,i} \neq 0$ and $T_{(1,i)}$ is invertible. Then we can find $\ell \in [n] \setminus \{1\}$ such that $t_{\ell,j} \neq 0$ and $T_{(1,i),(\ell,j)}$ is invertible. Let $T(a)$ be the matrix obtained by replacing the (ℓ, i) th entry in T with $a \in \mathbb{F}$. By the Laplace expansion of $T(a)_{(1,j)}$ with respect to the i th column, it follows that $\det(T(a)_{(1,j)}) = \lambda a + \gamma$ for some $\gamma \in \mathbb{F}$, where $\lambda = \det(T_{(1,j),(\ell,i)}) = \det(T_{(1,i),(\ell,j)}) \neq 0$. This means that $T(a)_{(1,j)}$ is invertible for all but one of $a \in \mathbb{F}$. Similarly, $\det(T(a)) = \lambda' a + \gamma'$ for some $\lambda' = \det(T_{(\ell,i)})$ and $\alpha' \in \mathbb{F}$,

which is not the zero polynomial. So for all but at most two of $a \in \mathbb{F}$, $\det(T(a))$ and $\det(T(a)_{(1,j)})$ are both non-zero. Since $|\mathbb{F}| > 2$, we can fix some $a \in \mathbb{F}$ such that $T(a)$ and $T(a)_{(1,j)}$ are both invertible. Note that the $(1, j)$ th entry of $T(a)$ is $t_{1,j} \neq 0$. If $T(a) \in \text{Conj}(\mathcal{S}_G)$, by Corollary 6.11, we can find an automorphism π of G sending 1 to j , which would contradict the assumption of orbits of $\text{Aut}(G)$. Recall that C'_i is the i th frontal slice of \mathbb{C}^{T^t} and denote by $C'_i(a)$ the i th frontal slice of $\mathbb{C}^{T(a)^t}$. We are going to use Lemma 6.8 to show $T(a) \in \text{Conj}(\mathcal{S}_G)$. To this end, we shall prove that, changing T to $T(a)$ does not affect $\text{colspan}(C'_k)$ and $\text{colspan}(TC_k)$ for any $k \in [n]$.

Firstly, we claim that $\text{colspan}(C'_k) = \text{colspan}(C'_k(a))$ for any $k \in [n]$. Since the only difference between T and $T(a)$ is the (ℓ, i) th entry, the claim clearly holds for all $k \in [n] \setminus \{\ell\}$. Furthermore, the only difference between C'_ℓ and $C'_\ell(a)$ as the linear combination of C_k 's is the coefficient of C_i , and the only problematic choice is when $a = 0$. To see this is not a problem, note that $t_{1,i}$ and $t_{1,j}$ are both non-zero, by (1) of Lemma 7.3, vertices i and j share the same out-neighborhood in G , thus $\text{colspan}(C_i) = \text{colspan}(C_j)$. Since $t_{\ell,j} \neq 0$, C_j is involved in both the linear combinations of C'_ℓ and $C'_\ell(a)$, which implies that $\text{colspan}(C'_\ell) = \text{colspan}(C'_\ell(a))$.

Secondly, we claim that $\text{colspan}(T(a)C_k) = \text{colspan}(TC_k)$ for any $k \in [n]$. To see this, note that $T(a)C_k$ and TC_k are different if and only if ℓ is in the out-neighborhood of k ; in other words, k is in the in-neighborhood of ℓ . Suppose this holds, and the only different column is t_ℓ^t and $t_\ell^t(a)$ (whose i th coordinate is replaced by a). Since $t_{1,j}$ and $t_{\ell,j}$ are both non-zero, Lemma 7.3 (2) ensures that the in-neighborhoods of 1 and ℓ are the same. It follows that 1 is also in the out-neighborhood of k , which means that t_1^t and t_ℓ^t are both in TC_k . Recall that the columns of C_k are standard column basis vectors and $C'_k = \sum_{k' \in [n]} t_{k,k'} C_{k'}$. By Observation 6.7, $\text{colspan}(C'_k)$ is spanned by the standard basis (or more precisely, the set $\{e_{k'}^t : t_{k,k'} \neq 0 \text{ \& } (k, k') \in E\}$). Since $\text{colspan}(TC_k) = \text{colspan}(C'_k)$ and the i th coordinate of t_1 is non-zero, $e_i^t \in \text{colspan}(C'_k)$. So changing the i th coordinate of t_ℓ to a , as long as the dimension of $\text{colspan}(T(a)C_k)$ does not decrease, does not affect $\text{colspan}(TC_k)$. These conclude that $\text{colspan}(TC_k) = \text{colspan}(T(a)C_k)$.

Since $T(a) \in \text{Conj}(\mathcal{S}_G)$, $T(a)_{(1,j)}$ are invertible, and the (i, j) th entry of $T(a)$ is $t_{i,j} \neq 0$, by Corollary 6.11, there exists an automorphism π of G sending 1 to j , which gives us the desired contradiction. \square

Claim 7.5 implies that all $T \in \text{Conj}(\mathcal{S}_G)$ are reducible with a common invariant subspace $\langle e_i, i \in S \rangle$, and thereby contradicts \mathcal{S}_G being conjugacy irreducible. \square

Proposition 7.6. *Let $G = ([n], E)$ be a directed graph, and $\mathcal{S}_G \leq M(n, \mathbb{F})$ be the matrix space associated with G . If \mathcal{S}_G is congruence irreducible, G is vertex transitive.*

Proof. We first claim that G must either have all-self-loops or no self-loops. If not, up to a permutation of vertices, we can assume that $E_{i,i} \notin \mathcal{S}_G$ but $E_{j,j} \in \mathcal{S}_G$ for all $i \in [s]$ and $j \in [n] \setminus [s]$. Note that for $T \in \text{Cong}(\mathcal{S}_G)$ and $j \in [n] \setminus [s]$, $TE_{j,j}T^t = [t_{k,j}t_{\ell,j}]_{k,\ell \in [n]} \in \mathcal{S}_G$. Since $E_{i,i} \notin \mathcal{S}_G$ for $i \in [s]$, $t_{i,j}t_{i,j} = 0$ for any $(i, j) \in [s] \times ([n] \setminus [s])$. This implies that $\text{Cong}(\mathcal{S}_G)$ is reducible with a common invariant subspace $\langle e_i \mid i \in [s] \rangle$, which is a contradiction.

Now we claim that, for $T \in \text{Cong}(\mathcal{S}_G)$, if $t_{i,j} \neq 0$ for some $i \neq j$, then there exists an automorphism π of G , such that $\pi(i) = j$. Since $T \in \text{Cong}(\mathcal{S}_G)$, by Corollary 6.4, there exists an automorphism σ of G such that $t_{i,\sigma(i)} \neq 0$ and $t_{\sigma^{-1}(j),j} \neq 0$. If $\sigma(i) = j$, then we are done. So assuming $\sigma(i) \neq j$ and $\sigma^{-1}(j) \neq i$, we embed $\sigma^{-1} \in \text{Aut}(G)$ as a matrix $P \in \text{GL}(n, \mathbb{F})$. Then we have $PT \in \text{Cong}(\mathcal{S}_G)$ with the $(\sigma(i), \sigma(i))$ th entry, the (j, j) th entry, and the $(\sigma(i), j)$ th entry all being non-zero. By Lemma 7.3, it implies that vertex $\sigma(i)$ and j share the same (in- and out-)neighborhood. Since G either has all-self-loops or no self-loops, the transposition of $\sigma(i)$ and j , denoted as σ' , is an automorphism of G . It follows that $\pi = \sigma' \circ \sigma$ is an automorphism of G sending i to j .

We prove the proposition by contradiction. Suppose \mathcal{S}_G is congruence irreducible but G is not vertex-transitive. Then $\text{Aut}(G)$ has at least two orbits. Let $S \subset [n]$ be one orbit of $\text{Aut}(G)$. Then for any $T \in \text{Cong}(\mathcal{S}_G)$ and for any $(i, j) \in S \times ([n] \setminus S)$, the (i, j) th entry of T must be zero, as otherwise there exists an automorphism of G sending i to j by the claim in the last paragraph, contradicting i and j are in different orbits of $\text{Aut}(G)$. It follows that $\text{Cong}(\mathcal{S}_G)$ is reducible due to the common invariant subspace $\langle e_i \mid i \in S \rangle$, contradicting the assumption that \mathcal{S}_G is congruence irreducible. \square

8 Connections to quantum information theory

In this section, we demonstrate connections of our results with some known results on quantum channels. We put such connections into two categories, which correspond to the two main approaches of viewing quantum channels as generalizations of graphs.

Before going into the details, we collect some basic notions from quantum information theory. To be consistent with the quantum information literature, we identify $u \in \mathbb{C}^n$ as *column vectors* and matrices $A \in M(n, \mathbb{C})$ acting on u from the left, i.e., $Au \in \mathbb{C}^n$. Recall that a linear map $\Phi : M(n, \mathbb{C}) \rightarrow M(n', \mathbb{C})$ is *completely positive* (CP) if for any $d \in \mathbb{N}$ and any positive semidefinite matrix $X \in M(d, \mathbb{C}) \otimes M(n, \mathbb{C})$, the matrix $(\text{id}_d \otimes \Phi)(X) \in M(d, \mathbb{C}) \otimes M(n', \mathbb{C})$ is positive semidefinite, where $\text{id}_d : M(d, \mathbb{C}) \rightarrow M(d, \mathbb{C})$ is the identity map on $M(d, \mathbb{C})$. Every CP linear map $\Phi : M(n, \mathbb{C}) \rightarrow M(n', \mathbb{C})$ admits the Choi–Kraus representation as follows. There exist matrices (Choi–Kraus operators) $E_1, \dots, E_m \in M(n' \times n, \mathbb{C})$ such that $\Phi(X) = \sum_{i=1}^m E_i X E_i^*$. While Choi–Kraus representations of a CP map Φ may not be unique, they span the same matrix space $\mathcal{S}_\Phi = \langle E_1, \dots, E_m \rangle$. We say Φ is *trace-preserving* (TP) if $\text{Tr}(\Phi(X)) = \text{Tr}(X)$ for any $X \in M(n, \mathbb{C})$, and Φ is *unital* if $\Phi(I_n) = I_{n'}$. In terms of Choi–Kraus operators, Φ is trace-preserving if $\sum_{i=1}^m E_i^* E_i = I_n$ and Φ is unital if $\sum_{i=1}^m E_i E_i^* = I_{n'}$. Quantum channels are trace-preserving CP maps, and unital quantum channels are natural quantum generalizations of doubly stochastic matrices.

8.1 Quantum channels and transition matrices

A transition matrix $P = (p_{i,j})_{i,j \in [n]} \in M(n, \mathbb{R}_{\geq 0})$ is a column stochastic matrix, i.e., it satisfies that $\sum_{i=1}^n p_{i,j} = 1$ for any $j \in [n]$. Such a matrix P describes a discrete Markov chain where the probability of moving from j to i is given by $p_{i,j}$.¹⁰ The underlying directed graph G of a transition matrix P is naturally defined with vertex set $[n]$ and arc set $E = \{(i, j) \mid p_{j,i} \neq 0\}$.

Quantum channels can be viewed as a natural quantum generalization of transition matrices. Indeed, associate each transition matrix $P = (p_{i,j})_{i,j \in [n]}$ with the CP map $\Phi_P : M(n, \mathbb{C}) \rightarrow M(n, \mathbb{C})$ defined as

$$\Phi_P(X) = \sum_{i,j=1}^n p_{i,j} E_{i,j} X E_{i,j}^* \quad (8.1)$$

for any $X \in M(n, \mathbb{C})$. It is clear that

$$\Phi_P(I_n) = \sum_{i,j=1}^n p_{i,j} E_{i,j}^* E_{i,j} = \sum_{i,j=1}^n p_{i,j} E_{j,j} = \sum_{j=1}^n \left(\sum_{i=1}^n p_{i,j} \right) E_{j,j} = \sum_{j=1}^n E_{j,j} = I_n,$$

¹⁰In the theory of discrete Markov chains, the transition matrix normally acts on probability row vectors from right. For consistency with quantum channels, we translate the action to acting from the left on probability column vectors; see also [BBLM20].

where the second last equality uses the fact that P is column stochastic. Thus Φ_P is a quantum channel. To realize the action of P on probability vectors, we apply Φ_P on diagonal density matrices (positive semidefinite matrices with unit trace) [BBLM20]. Observe that $\mathcal{S}_{\Phi_P} = \langle \sqrt{p_{i,j}} E_{i,j} \mid p_{i,j} \neq 0 \rangle = \mathcal{S}_G^*$, where G is the underlying directed graph of P and the adjoint space \mathcal{S}^* of \mathcal{S} is defined as $\{B^* : B \in \mathcal{S}\}$. This illustrates a natural way to generalize directed graphs to quantum channels.

Irreducibility of transition matrices. We first connect the irreducibility of a quantum channel $\Phi : M(n, \mathbb{C}) \rightarrow M(n, \mathbb{C})$ with the irreducibility of its underlying matrix space \mathcal{S}_Φ . A quantum channel Φ is *irreducible*, if the only orthogonal projections P (i.e., $P^2 = P$ and $P^* = P$) satisfying $\Phi(PM(n, \mathbb{C})P) \leq PM(n, \mathbb{C})P$ are 0 and I_n (cf. [EHK78, Wol12]). We first observe that the irreducibility notion for quantum channels coincides with the irreducibility notion for their underlying matrix spaces.

Proposition 8.1. *A CP map $\Phi : M(n, \mathbb{C}) \rightarrow M(n, \mathbb{C})$ is irreducible if and only if \mathcal{S}_Φ is irreducible.*

Proof. Let E_1, \dots, E_m be a set of Choi–Kraus operators for Φ . We first show that Φ being reducible implies that \mathcal{S}_Φ is reducible. Let P be a non-trivial projection such that $\Phi(PM(n, \mathbb{C})P) \leq PM(n, \mathbb{C})P$. Let $u_1, \dots, u_k \in \mathbb{C}^n$ be an orthonormal basis of $\text{colspan}(P)$, so $PM(n, \mathbb{C})P = PM(n, \mathbb{C})P^* = \langle u_i u_j^* \mid i, j \in [k] \rangle$. Then for each $j \in [k]$,

$$\Phi(u_j u_j^*) = \sum_{i=1}^m E_i u_j u_j^* E_i^* \in \langle u_i u_j^* \mid i, j \in [k] \rangle.$$

We now claim that $\text{colspan}(\Phi(u_j u_j^*)) = \langle E_i u_j \mid i \in [m] \rangle$. First, it is straightforward to verify that $\text{colspan}(\Phi(u_j u_j^*)) \leq \langle E_i u_j \mid i \in [m] \rangle$. Second, for any $v \in \text{colspan}(\Phi(u_j u_j^*))^\perp$, we have $v^* \Phi(u_j u_j^*) v = 0$. Expanding the left-hand side, we have

$$0 = v^* \Phi(u_j u_j^*) v = \sum_{i=1}^m v^* E_i u_j u_j^* E_i^* v = \sum_{i=1}^m |v^* E_i u_j|^2.$$

Thus $v \in \langle E_i u_j \mid i \in [m] \rangle^\perp$, showing that $\text{colspan}(\Phi(u_j u_j^*))^\perp \leq \langle E_i u_j \mid i \in [m] \rangle^\perp$. The claim is then proved.

It follows that $\text{colspan}(\Phi(u_j u_j^*)) = \langle E_i u_j \mid i \in [m] \rangle \leq \text{colspan}(P)$ for every $j \in [k]$. This implies that $\text{colspan}(P)$ is a non-trivial invariant subspace of \mathcal{S}_Φ , thus \mathcal{S}_Φ is reducible.

We now show that \mathcal{S}_Φ being reducible implies that Φ is reducible. Let U be a non-trivial invariant subspace of \mathcal{S}_Φ of dimension d and let P be the orthogonal projection on U . It is not hard to verify that $\Phi(PM(n, \mathbb{C})P) \leq PM(n, \mathbb{C})P$. \square

Recall that a transition matrix P is irreducible if its underlying directed graph G is strongly connected. Thus, Theorem 1.12 formally bridges the irreducibility of transition matrices and quantum channels.

Corollary 8.2. *Let P be a transition matrix, G be its underlying directed graph and Φ_P be the associated quantum channel given in Eq. (8.1). Then P is irreducible if and only if Φ_P is irreducible.*

Proof. Let $G^t = ([n], E^t)$ be the transpose graph of $G = ([n], E)$, where $E^t = \{(i, j) \mid (j, i) \in E\}$. Then $\mathcal{S}_G^* = \mathcal{S}_{G^t}$. Moreover, G is strongly connected if and only if G^t is strongly connected. The equivalence then follows from Theorem 1.12 and $\mathcal{S}_{\Phi_P} = \mathcal{S}_G^* = \mathcal{S}_{G^t}$. \square

Vertex transitivity. We then turn to vertex transitivity. Since every vertex-transitive graph G is regular, its normalized adjacency matrix A is a symmetric transition matrix. We can associate to G the following quantum channel Φ_G in the spirit of Eq. (8.1):

$$\Phi_G(X) = \frac{1}{d} \sum_{(i,j) \in E} E_{i,j} X E_{i,j}^* \quad (8.2)$$

for every $X \in M(n, \mathbb{C})$, where d denotes the in-degree of any vertex in G . It is clear that $\mathcal{S}_{\Phi_G} = \mathcal{S}_G$.

An interesting result in [BBLM20] connects vertex-transitive graphs with irreducibly covariant channels. A quantum channel $\Phi : M(n, \mathbb{C}) \rightarrow M(n, \mathbb{C})$ is *irreducibly covariant*, if there exists a compact group Γ and a continuous irreducible unitary representation $U : \Gamma \rightarrow U(n)$ such that for any $g \in \Gamma$ and $X \in M(n, \mathbb{C})$, we have $\Phi(U(g)XU(g)^*) = U(g)\Phi(X)U(g)^*$.

Proposition 8.3 ([BBLM20, Proposition 3.8]). *G is vertex-transitive if and only if Φ_G (as defined in Eq. (8.2)) is irreducibly covariant.*

We can deduce Proposition 8.3 easily from Theorem 1.16 and Proposition 7.1.

Proof. For the if direction, let $U : \Gamma \rightarrow U(n)$ be the irreducible unitary representation of some compact group Γ such that for any $g \in \Gamma$ and $X \in M(n, \mathbb{C})$, we have $\Phi(U(g)XU(g)^*) = U(g)\Phi(X)U(g)^*$. In this case, for every $g \in \Gamma$, the map $\Phi' : M(n, \mathbb{C}) \rightarrow M(n, \mathbb{C})$ given by $\Phi'(X) = \Phi(U(g)XU(g)^*)$ and the map $\Phi'' : M(n, \mathbb{C}) \rightarrow M(n, \mathbb{C})$ given by $\Phi''(X) = U(g)\Phi(X)U(g)^*$ are equivalent. This implies that $\mathcal{S}_{\Phi_G}U(g) = \mathcal{S}_{\Phi'} = \mathcal{S}_{\Phi''} = U(g)\mathcal{S}_{\Phi_G}$. It follows that $U(\Gamma) \leq \text{Conj}(\mathcal{S}_G)$, and, since $U(\Gamma)$ is irreducible, $\text{Conj}(\mathcal{S}_G)$ is irreducible. By Theorem 1.16, we can conclude that G is vertex-transitive.

For the only if direction, since G is vertex-transitive, Proposition 7.1 proves that the matrix group H generated by the automorphism group of G (embedded as a subgroup of $\text{GL}(n, \mathbb{C})$) and the group of diagonal unitary matrices is irreducible. It is straightforward to verify that H is compact and $\Phi(UXU^*) = U\Phi(X)U^*$ for any $U \in H$. \square

Summarizing Theorem 1.16 and Proposition 8.3, we have the following.

Corollary 8.4. *The following are equivalent:*

- (1) G is vertex-transitive;
- (2) Φ_G is irreducibly covariant;
- (3) \mathcal{S}_G is conjugacy irreducible;
- (4) \mathcal{S}_G is congruence irreducible.

8.2 Quantum expanders and dimension expanders.

Expander graphs are an important family of graphs which are sparse while highly connected. A natural expansion property for matrix spaces is the spectral expansion, which is mostly studied for *quantum expanders* [Has07, Har08, BASTS10, Pis14]. Recall that the spectral expansion of a d -regular graph G is the second largest absolute value of the eigenvalues of its adjacency matrix A_G . On the other hand, from a d -regular graph $G = ([n], E)$, we have the linear map $\Phi_G : M(n, \mathbb{C}) \rightarrow M(n, \mathbb{C})$ (as given in Eq. (8.2)), such that \mathcal{S}_G is the matrix space spanned by the Choi-Kraus operator of Φ_G . Define the spectral expansion of a linear map $\Phi : M(n, \mathbb{C}) \rightarrow M(n, \mathbb{C})$ as the second largest absolute value of its eigenvalues.

Theorem 8.5 ([BBLM20, Proposition 3.7]). *For any d -regular graph G , the spectral expansion of G equals the spectral expansion of Φ_G (as given in Eq. (8.2)).*

Another type of expansion property for matrix spaces is dimension expansion. We say a matrix space $\mathcal{S} \leq M(n, \mathbb{F})$ is a degree- d *dimension expander* if $\dim(\mathcal{S}) = d$ and there exists a constant α such that for any subspace $V \leq \mathbb{F}^n$ with $\dim(V) \leq n/2$, we have $\dim(\mathcal{S}(V)) \geq (1 + \alpha) \dim(V)$.

In a follow-up paper [LQW⁺22], we discuss relationships between the different linear-algebraic notions of expansion. We also prove extensions of Theorem 8.5 to other notions of expansion; these theorems, like Theorem 8.5 can be viewed as additional inherited correspondences. Indeed, they state that a certain graph-theoretic optimum is equal to an associated linear-algebraic optimum.

Dimension expanders have been shown to connect with *monotone expanders* [DS09, DW10]. Let \mathcal{F} be a set of d (partial) monotone functions $f_1, \dots, f_d : [n] \rightarrow [n]$. Let $E = \{(i, f_j(i)) \mid i \in [n], j \in [d]\}$. We call $G_{\mathcal{F}} = ([n] \times [n], E)$ a *degree- d monotone graph* with respect to \mathcal{F} . A degree- d monotone graph is a degree- d monotone expander if there exists a constant α such that every set $A \subseteq [n]$ of left vertices of size at most $n/2$ has at least $(1 + \alpha)|A|$ neighbors. Given a degree- d monotone graph $G_{\mathcal{F}} = ([n], E)$, construct a matrix space $\mathcal{S}_{\mathcal{F}} = \langle \sum_{i \in [n]} e_i e_{f_j(i)}^t \mid j \in [d] \rangle$.

Theorem 8.6 ([DS09, DW10]). *For a set of d monotone functions $\mathcal{F} = \{f_1, \dots, f_d\}$, $G_{\mathcal{F}}$ is a degree- d monotone expander with constant α if and only if the matrix space $\mathcal{S}_{\mathcal{F}}$ is a degree- d dimension expander with constant α .*

Note that the matrix space $\mathcal{S}_{\mathcal{F}}$ is different from the graphical matrix space $\mathcal{S}_{G_{\mathcal{F}}}$. Combined with Bourgain’s explicit construction of constant-degree monotone expanders [Bou09, BY13], one directly obtains an explicit construction of constant-degree dimension expanders by Theorem 8.6. Note that over fields \mathbb{F} of characteristic 0, one may use other expanders to construct dimension expanders [LZ08, LQW⁺22].

8.3 Quantum channels and zero-error communication

Another way to connect graphs with quantum channels is via a quantum version of the zero-error capacity problem. In Shannon’s seminal work [Sha56], an *undirected* graph is associated to each classical communication channel, whose the independence number equals the largest number of zero-error messages one can send through the channel. Replacing the classical channel by a quantum channel Φ , the role of graphs is taken by the *operator system* associated with Φ [DSW13]. More precisely, let $E_1, \dots, E_m \in M(n' \times n, \mathbb{C})$ be a set of Choi–Kraus operators of $\Phi : M(n, \mathbb{C}) \rightarrow M(n', \mathbb{C})$. The operator system of Φ is $\mathcal{F}_{\Phi} = \langle E_i^* E_j \mid i, j \in [m] \rangle \leq M(n, \mathbb{C})$, which is self-adjoint (i.e., $X \in \mathcal{F}_{\Phi}$ implies $X^* \in \mathcal{F}_{\Phi}$) and $I_n \in \mathcal{F}$ (since Φ is trace-preserving). On the other hand, for every self-adjoint $\mathcal{F} \leq M(n, \mathbb{C})$ containing I_n , one can find a quantum channel Φ such that $\mathcal{F}_{\Phi} = \mathcal{F}$. The authors of [DSW13] viewed operator systems as a non-commutative generalization of graphs, so called them non-commutative graphs.

The correspondence between (undirected) graphs and operator systems is the following. Let $G = ([n], E)$ be an undirected graph. We then define its corresponding operator system by

$$\mathcal{F}_G = \langle E_{i,j} \mid \{i, j\} \in E \text{ or } i = j \in [n] \rangle. \quad (8.3)$$

Note that \mathcal{F}_G and the graphical matrix space \mathcal{S}_G are different. We identify each undirected graph $G = ([n], E)$ as a directed graph $\hat{G} = ([n], \hat{E})$, where $\hat{E} = \{(i, j) \mid \{i, j\} \in E \text{ or } i = j \in [n]\}$. In other words, we obtain \hat{G} by identifying each undirected edge as two directed edges and adding self-loops to each vertex of G . Through this identification, we have $\mathcal{F}_G = \mathcal{S}_{\hat{G}}$.

Connectivity of operator systems. The connectivity of operator systems was studied in [CDS21]. For an operator system $\mathcal{F} \leq M(n, \mathbb{C})$, we say \mathcal{F} is connected, if there is no non-trivial orthogonal projection $P \in M(n, \mathbb{C})$, such that $PS(I_n - P) = \{0\}$ (cf. [CDS21, Theorem 3.3]).

Corollary 8.7. *Let $G = ([n], E)$ be an undirected graph and $\hat{G} = ([n], \hat{E})$ be the directed graph obtained as above. Then the following are equivalent:*

- (1) G is connected;
- (2) \hat{G} is strongly connected;
- (3) \mathcal{F}_G is connected;
- (4) $\mathcal{S}_{\hat{G}}$ is irreducible.

Proof. The equivalence between (1) and (2) is direct to verify. The equivalence between (1) and (3) is shown in [CDS21, Corollary 3.4] and the equivalence between (2) and (4) is due to Theorem 1.12. Note that the connectivity of an operator system \mathcal{F} exactly implies that \mathcal{F} is irreducible. Thus (3) and (4) are equivalent. \square

Isomorphic operator systems. The natural notion of equivalence between operator systems is the so-called *unital and complete order isomorphism*. Namely, two operator systems \mathcal{F}_1 and \mathcal{F}_2 are unital, complete order isomorphic if there is a unital linear map $\phi : \mathcal{F}_1 \rightarrow \mathcal{F}_2$ which is one-to-one, onto and completely positive for both ϕ and ϕ^{-1} . For undirected graphs G and H , the equivalence between \mathcal{F}_G and \mathcal{F}_H is compatible with the isomorphism of G and H , and admits a simple form:

Theorem 8.8 ([OP15, Theorem 3.3]). *Let G and H be undirected graphs. Then the following are equivalent:*

- (1) G and H are isomorphic;
- (2) \mathcal{F}_G and \mathcal{F}_H are unital and complete order isomorphic;
- (3) There exists a unitary U such that $U\mathcal{F}_GU^* = \mathcal{F}_H$.

We note that Proposition 6.2 could be viewed as a generalization of the equivalence between (1) and (3) in Theorem 8.8, as shown in the following result.

Corollary 8.9. *Let G and H be undirected graphs. Then the following are equivalent:*

- (1) G and H are isomorphic;
- (2) There exists an invertible matrix T such that $T\mathcal{F}_GT^* = \mathcal{F}_H$;

Proof. Let \hat{G} and \hat{H} be the corresponding directed graphs obtained as above. Note that G and H are isomorphic if and only if \hat{G} and \hat{H} are isomorphic. Then the equivalence between (1) and (2) follows from Proposition 6.2 and $\mathcal{F}_G = \mathcal{S}_{\hat{G}}$ and $\mathcal{F}_H = \mathcal{S}_{\hat{H}}$. \square

In Theorem 8.8, to show the equivalence between (1) and (3), the more difficult direction is to show (3) \Rightarrow (1). This is immediate from our Corollary 8.9.

Acknowledgments

Y.Q. would like to thank George Glauberger for helping him with Proposition 7.1. We would also like to thank an anonymous referee for helpful comments and intriguing suggestions.

References

- [Ada62] J. F. Adams. Vector fields on spheres. *Ann. of Math. (2)*, 75:603–632, 1962.
- [Alb50] Abraham A. Albert. A theory of power-associative commutative algebras. *Trans. Amer. Math. Soc.*, 69:503–527, 1950.
- [ALP65] J. F. Adams, Peter D. Lax, and Ralph S. Phillips. On matrices whose real linear combinations are nonsingular. *Proc. Amer. Math. Soc.*, 16(2):318–322, 1965.
- [Atk80] M. D. Atkinson. Spaces of matrices with several zero eigenvalues. *Bull. Lond. Math. Soc.*, 12(2):89–95, 1980.
- [Bae38] Reinhold Baer. Groups with abelian central quotient group. *Trans. Amer. Math. Soc.*, 44(3):357–386, 1938.
- [BASTS10] Avraham Ben-Aroya, Oded Schwartz, and Amnon Ta-Shma. Quantum expanders: Motivation and construction. *Theory Comput.*, 6(3):47–79, 2010.
- [BBLM20] Tom Bannink, Jop Briët, Farrokh Labib, and Hans Maassen. Quasirandom quantum channels. *Quantum*, 4:298, 2020.
- [BCG⁺21] Xiaohui Bei, Shiteng Chen, Ji Guan, Youming Qiao, and Xiaoming Sun. From independent sets and vertex colorings to isotropic spaces and isotropic decompositions: Another bridge between graphs and alternating matrix spaces. *SIAM J. Comput.*, 50(3):924–971, 2021.
- [Ber84] Stuart J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Inform. Process. Lett.*, 18(3):147–150, 1984.
- [BFS99] Jonathan F. Buss, Gudmund S. Frandsen, and Jeffrey O. Shallit. The computational complexity of some problems of linear algebra. *J. Comput. Syst. Sci.*, 58(3):572–596, 1999.
- [BJG08] Jørgen Bang-Jensen and Gregory Z. Gutin. *Digraphs: theory, algorithms and applications*. Springer, 2008.
- [Bou09] Jean Bourgain. Expanders and dimensional expansion. *C. R. Math. Acad. Sci. Paris*, 347(7-8):357–362, 2009.
- [BS20] Jendrik Brachter and Pascal Schweitzer. On the Weisfeiler-Leman dimension of finite groups. In *LICS '20: 35th Annual ACM/IEEE Symposium on Logic in Computer Science, Saarbrücken, Germany, July 8-11, 2020*, pages 287–300. ACM, 2020.
- [BY13] Jean Bourgain and Amir Yehudayoff. Expansion in $SL_2(\mathbb{R})$ and monotone expanders. *Geom. Funct. Anal.*, 23(1):1–41, 2013.
- [CDS21] Javier Alejandro Chavez-Domnguez and Andrew T. Swift. Connectivity for quantum graphs. *Linear Algebra Appl.*, 608:37–53, 2021.
- [Coh75] P. M. Cohn. The word problem for free fields: A correction and an addendum. *J. Symbolic Logic*, 40(1):69–74, 1975.

- [Die48] Jean Dieudonné. Sur une généralisation du groupe orthogonal à quatre variables. *Arch. Math.*, 1(4):282–287, 1948.
- [Die17] Reinhard Diestel. *Graph Theory*, volume 173 of *Springer Graduate Texts in Mathematics*. Springer, 5th edition, 2017.
- [DS09] Zeev Dvir and Amir Shpilka. Towards dimension expanders over finite fields. *Combinatorica*, 31(3):305320, 2009.
- [dSP13] Clément de Seguins Pazzis. On Gerstenhabers theorem for spaces of nilpotent matrices over a skew field. *Linear Algebra Appl.*, 438(11):4426–4438, 2013.
- [DSW13] Runyao Duan, Simone Severini, and Andreas Winter. Zero-error communication via quantum channels, noncommutative graphs, and a quantum Lovász number. *IEEE Trans. Inform. Theory*, 59(2):1164–1174, 2013.
- [DW10] Zeev Dvir and Avi Wigderson. Monotone expanders: Constructions and applications. *Theory Comput.*, 6(12):291–308, 2010.
- [Edm67] Jack Edmonds. Systems of distinct representatives and linear algebra. *J. Res. Nat. Bur. Standards Sect. B*, 71B:241–245, 1967.
- [EH88] David Eisenbud and Joe Harris. Vector spaces of matrices of low rank. *Adv. Math.*, 70(2):135–155, 1988.
- [EHK78] David E. Evans and Raphael Høegh-Krohn. Spectral properties of positive maps on C^* -algebras. *J. Lond. Math. Soc.*, 2(2):345–355, 1978.
- [FGS19] Vyacheslav Futorny, Joshua A. Grochow, and Vladimir V. Sergeichuk. Wildness for tensors. *Linear Algebra Appl.*, 566:212–244, 2019.
- [FGT21] Stephen Fenner, Rohit Gurjar, and Thomas Thierauf. Bipartite perfect matching is in Quasi-NC. *SIAM J. Comput.*, 50(3):218–235, 2021.
- [Fla62] H. Flanders. On spaces of linear transformations with bounded rank. *J. Lond. Math. Soc.*, 1(1):10–16, 1962.
- [Fro97] F. G. Frobenius. Über die Darstellung der endlichen Gruppen durch lineare substitutionen. *Stiz. Kön. Preuss. Akad. Wiss. Berlin*, pages 944–1015, 1897.
- [FS13] Michael A. Forbes and Amir Shpilka. Explicit Noether normalization for simultaneous conjugation via polynomial identity testing. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 16th International Workshop, APPROX 2013, and 17th International Workshop, RANDOM 2013*, volume 8096 of *Lecture Notes in Computer Science*, pages 527–542. Springer, 2013.
- [Ger58] Murray Gerstenhaber. On nilalgebras and linear varieties of nilpotent matrices, I. *Amer. J. Math.*, 80(3):614–622, 1958.
- [GGOW20] Ankit Garg, Leonid Gurvits, Rafael M. Oliveira, and Avi Wigderson. Operator scaling: Theory and applications. *Found. Comput. Math.*, 20(2):223–290, 2020.
- [Gur04] Leonid Gurvits. Classical complexity and quantum entanglement. *J. Comput. Syst. Sci.*, 69(3):448–484, 2004.

- [Hal35] P. Hall. On representatives of subsets. *J. Lond. Math. Soc.*, 1(1):26–30, 1935.
- [Har08] Aram W. Harrow. Quantum expanders from any classical Cayley graph expander. *Quantum Inf. Comput.*, 8(8):715–721, 2008.
- [Has07] Matthew B. Hastings. Random unitaries give quantum expanders. *Phys. Rev. A*, 76:032315, 2007.
- [HH21] Masaki Hamada and Hiroshi Hirai. Computing the nc-rank via discrete convex optimization on CAT (0) spaces. *SIAM J. Appl. Algebra Geom.*, 5(3):455–478, 2021.
- [HJ85] Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.
- [HQ21] Xiaoyu He and Youming Qiao. On the Baer-Lovász-Tutte construction of groups from graphs: Isomorphism types and homomorphism notions. *European J. Combin.*, 98:Paper No. 103404, 12pp, 2021.
- [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.
- [IKS10] Gábor Ivanyos, Marek Karpinski, and Nitin Saxena. Deterministic polynomial time algorithms for matrix completion problems. *SIAM J. Comput.*, 39(8):3736–3751, 2010.
- [IQS17] Gábor Ivanyos, Youming Qiao, and K. V. Subrahmanyam. Non-commutative Edmonds’ problem and matrix semi-invariants. *Comput. 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. *Comput. Complexity*, 27(4):561–593, 2018.
- [Kar72] Richard M. Karp. Reducibility among combinatorial problems. In *Complexity of computer computations (Proc. Sympos., IBM Thomas J. Watson Res. Center, Yorktown Heights, N.Y., 1972)*, pages 85–103, 1972.
- [KI04] Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Comput. Complexity*, 13(1/2):1–46, 2004.
- [KUW86] Richard M. Karp, Eli Upfal, and Avi Wigderson. Constructing a perfect matching is in random NC. *Combinatorica*, 6(1):35–48, 1986.
- [Lan05] S. Lang. *Algebra*. Graduate Texts in Mathematics. Springer New York, 2005.
- [LLW88] N. Linial, L. Lovász, and A. Wigderson. Rubber bands, convex embeddings and graph connectivity. *Combinatorica*, 8(1):91–102, 1988.
- [Lov79] László Lovász. On determinants, matchings, and random algorithms. In *Fundamentals of computation theory (Proc. Conf. Algebraic, Arith. and Categorical Methods in Comput. Theory, Berlin/Wendisch-Rietz, 1979)*, volume 2 of *Math. Res.*, pages 565–574. Akademie-Verlag, Berlin, 1979.
- [Lov89] László Lovász. Singular spaces of matrices and their application in combinatorics. *Bol. Soc. Brasil. Mat. (N.S.)*, 20(1):87–99, 1989.

- [LQ17] Yinan Li and Youming Qiao. Linear algebraic analogues of the graph isomorphism problem and the Erdős–Rényi model. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 463–474. IEEE Computer Society, 2017.
- [LQ20] Yinan Li and Youming Qiao. Group-theoretic generalisations of vertex and edge connectivities. *Proc. Amer. Math. Soc.*, 148(11):4679–4693, 2020.
- [LQW⁺22] Yinan Li, Youming Qiao, Avi Wigderson, Yuval Wigderson, and Chuanqi Zhang. On linear-algebraic notions of expansion, 2022. arXiv:2212.13154.
- [LZ08] Alexander Lubotzky and Efim Zelmanov. Dimension expanders. *J. Algebra*, 319(2):730–738, 2008.
- [Mes85] Roy Meshulam. On the maximal rank in a subspace of matrices. *Q. J. Math.*, 36(2):225–229, 1985.
- [MMS12] G. W. MacDonald, J. A. MacDougall, and L. G. Sweet. On the dimension of linear spaces of nilpotent matrices. *Linear Algebra Appl.*, 436(7):2210–2230, 2012.
- [MOR91] Ben Mathes, Matjaž Omladič, and Heydar Radjavi. Linear spaces of nilpotent matrices. *Linear Algebra Appl.*, 149:215–225, 1991.
- [MVV87] Ketan Mulmuley, Umesh V. Vazirani, and Vijay V. Vazirani. Matching is as easy as matrix inversion. *Combinatorica*, 7(1):105–113, 1987.
- [MW21] Visu Makam and Avi Wigderson. Singular tuples of matrices is not a null cone (and the symmetries of algebraic varieties). *J. Reine Angew. Math.*, 780:79–131, 2021.
- [OP15] Carlos M. Ortiz and Vern I. Paulsen. Lovász theta type norms and operator systems. *Linear Algebra Appl.*, 477:128–147, 2015.
- [Pis14] Gilles Pisier. Quantum expanders and geometry of operator spaces. *J. Eur. Math. Soc.*, 16:11831219, 2014.
- [Qia20] Youming Qiao. Turán and Ramsey problems for alternating multilinear maps, 2020. arXiv:2007.12820.
- [Qia21] Youming Qiao. Enumerating alternating matrix spaces over finite fields with explicit coordinates. *Discrete Math.*, 344(11):112580, 2021.
- [Ros22] Tobias Rossmann. Enumerating conjugacy classes of graphical groups over finite fields. *Bull. Lond. Math. Soc.*, 2022. To appear.
- [RV22] Tobias Rossmann and Christopher Voll. Groups, graphs, and hypergraphs: Average sizes of kernels of generic matrices with support constraints. *Mem. Amer. Math. Soc.*, 2022. To appear.
- [RW19] Orit E. Raz and Avi Wigderson. Subspace arrangements, graph rigidity and derandomization through submodular optimization. In *Building Bridges. II*, volume 28 of *Bolyai Soc. Math. Stud.*, pages 377–415. Springer, Berlin, 2019.
- [Sch80] J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. Assoc. Comput. Mach.*, 27(4):701–717, 1980.

- [Ser85] V. N. Serežkin. On linear transformations preserving nilpotency. *Izv. Akad. Nauk BSSR, Ser. Fiz.-Mat. Nauk*, 6:46–50, 1985.
- [Sha56] Claude E. Shannon. The zero error capacity of a noisy channel. *Institute of Radio Engineers Transactions on Information Theory*, IT-2:8–19, 1956.
- [ST17] Ola Svensson and Jakub Tarnawski. The matching problem in general graphs is in quasi-NC. In *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.
- [SY09] Amir Shpilka and Amir Yehudayoff. Arithmetic circuits: a survey of recent results and open questions. *Found. Trends Theor. Comput. Sci.*, 5(3-4):207–388, 2009.
- [Tut47] W. T. Tutte. The factorization of linear graphs. *J. London Math. Soc.*, 22:107–111, 1947.
- [Val79] Leslie G. Valiant. Completeness classes in algebra. In *Proceedings of the 11h Annual ACM Symposium on Theory of Computing, April 30-May 2, 1979, Atlanta, Georgia, USA*, pages 249–261. ACM, 1979.
- [VF17] E. O. Quintero Vanegas and Juan C. Gutierrez Fernandez. Nilpotent linear spaces and Albert’s problem. *Linear Algebra Appl.*, 518:57–78, 2017.
- [Wea21] Nik Weaver. Quantum graphs as quantum relations. *J. Geom. Anal.*, 31:1–23, 2021.
- [Wol12] Michael M. Wolf. Quantum channels & operations: Guided tour. Lecture Notes, 2012.
- [Zha06] Fuzhen Zhang. *The Schur complement and its applications*, volume 4. Springer, 2006.
- [Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In *Symbolic and Algebraic Computation, EUROSAM ’79, An International Symposium on Symbolic and Algebraic Computation, Marseille, France, June 1979*, volume 72 of *Lecture Notes in Computer Science*, pages 216–226. Springer, 1979.
- [ZRP⁺09] R. Zenklusen, B. Ries, C. Picouleau, D. de Werra, M.-C. Costa, and C. Bentz. Blockers and transversals. *Discrete Math.*, 309(13):4306–4314, 2009.

A On finding cycles in directed graphs

Proposition A.1. *There is an NC-reduction from finding cycles in directed graphs to finding perfect matchings in bipartite graphs.*

Proof. Let $G = ([n], E)$ be a directed graph. We construct a bipartite graph $H = (L \times R, F)$ as follows: Let L and R be two copies of $[n]$. To distinguish them, denote the vertices of L as $i_{L,1}, i_{L,2}$ for all $i \in [n]$ and the vertices of R as $j_{R,1}, j_{R,2}$ for all $j \in [n]$. The edge set F is constructed as follows: For every $(i, j) \in E$, set $(i_{L,1}, j_{R,1}) \in F$. For every $i \in [n]$, let $(i_{L,1}, i_{R,2}), (i_{L,2}, i_{R,2}), (i_{L,2}, i_{R,1}) \in F$. The second set of edges is exactly length-3 paths from $i_{L,1}$ to $i_{R,1}$, which will be denoted by P_i for each $i \in [n]$.

Consider perfect matchings in H , which must take edges from each P_i . For a fixed perfect matching M , divide $[n]$ into two disjoint sets S and T , where S contains those i where the edge $(i_{L,2}, i_{R,2}) \in M$, and T contains those i where the edges $(i_{L,1}, i_{R,2}), (i_{L,2}, i_{R,1}) \in M$. Note that if

$S \neq \emptyset$, the edges in M which start from $i_{L,1}$ for $i \in S$ will end in some $j_{R,1}$ for $j \in S$, and these edges form a cycle or disjoint union of cycles in G .

The only problem is S can be empty. To resolve this issue, one can remove the edges in P_1 from F and repeat the above steps, which will yield cycles in G containing 1 (if there is one). Repeating this for all i yields the desired reduction. \square

Proposition A.1 implies that finding a cycle on a directed graph can be done in quasi-NC by [FGT21]. However, it is not hard to show that finding a cycle in a directed graph is in NC.

Proposition A.2. *Finding a cycle in a directed graph can be solved in NC.*

Proof sketch. Let $G = ([n], E)$ be a directed graph. Let A_G be the adjacency matrix of G . For each entry of A_G we store a walk of length 1. Then for $k \in [\log n]$, compute $A_G^{2^k}$ recursively using $A_G^{2^{k-1}}$. Furthermore, for each entry of $A_G^{2^k}$, if it is non-zero, store a walk of length 2^k (by concatenating two walks). In this way, if $A_G^{2^{\lceil \log n \rceil}}$ is non-zero (which happens if and only if G is cyclic), we obtain a walk of length $\geq n$, from which we can extract a cycle easily. All the above steps run in NC². \square

Proposition A.3. *Matrix space nil testing reduces to symbolic determinant identity testing.*

Proof. Let $\mathcal{S} = \langle B_1, \dots, B_m \rangle \leq M(n, \mathbb{F})$ be a matrix space. Let x_1, \dots, x_m be variables, and construct a symbolic matrix $\mathbf{B} := x_1 B_1 + \dots + x_m B_m$. Note that \mathcal{S} is nil if and only if \mathbf{B}^n is the all-zero matrix. By [FS13], testing if a particular entry of \mathbf{B}^n is the zero polynomial reduces to SDIT. So we can use n^2 instances of SDIT to solve the matrix space nil testing problem. \square