

*Flanders' theorem for many matrices under
commutativity assumptions*

De Teran, Fernando and Lippert,
Ross and Nakatsukasa, Yuji and Noferini, Vanni

2013

MIMS EPrint: **2013.20**

Manchester Institute for Mathematical Sciences
School of Mathematics

The University of Manchester

Reports available from: <http://eprints.maths.manchester.ac.uk/>

And by contacting: The MIMS Secretary
School of Mathematics
The University of Manchester
Manchester, M13 9PL, UK

ISSN 1749-9097

Flanders' theorem for many matrices under commutativity assumptions

Fernando De Terán¹

Universidad Carlos III de Madrid, Avda. Universidad 30, 28911, Leganés, Spain

Ross A. Lippert

123 W. 92 Street #1, New York, NY 10025, USA

Yuji Nakatsukasa²

Department of Mathematical Informatics, Graduate School of Information Science and Technology, University of Tokyo, Tokyo 113-8656, Japan

Vanni Noferini³

School of Mathematics, The University of Manchester, Manchester, M13 9PL, UK

Abstract

We analyze the relationship between the Jordan canonical form of products, in different orders, of k square matrices A_1, \dots, A_k . Our results extend some classical results by H. Flanders. Motivated by a generalization of Fiedler matrices, we study permuted products of A_1, \dots, A_k under the assumption that the graph of non-commutativity relations of A_1, \dots, A_k is a forest. Under this condition, we show that the Jordan structure of all nonzero eigenvalues is the same for all permuted products. For the eigenvalue zero, we obtain an upper bound on the difference between the sizes of Jordan blocks for any two permuted products, and we show that this bound is attainable. For $k = 3$ we show that, moreover, the bound is exhaustive.

AMS subject classification: 05A05, 05C50, 15A18, 15A21.

Keywords: eigenvalue, Jordan canonical form, Segré characteristic, product of matrices, permuted products, Flanders' theorem, forest, cut-flip.

Email addresses: fteran@math.uc3m.es (Fernando De Terán), ross.lippert@gmail.com (Ross A. Lippert), nakatsukasa@mist.i.u-tokyo.ac.jp (Yuji Nakatsukasa), vanni.noferini@manchester.ac.uk (Vanni Noferini)

¹Supported by the Ministerio de Economía y Competitividad of Spain through grants MTM-2009-09281 and MTM-2012-32542.

²Supported in part by EPSRC grant EP/I005293/1.

³Supported by ERC Advanced Grant MATFUN-267526.

1. Introduction

The *Jordan canonical form* (JCF) is a familiar canonical form under similarity of square matrices. It consists of a direct sum of *Jordan blocks* associated with eigenvalues, and it is unique up to permutation of these blocks [8, §3.1]. We assume throughout the paper that, for a given eigenvalue λ , the Jordan blocks at λ in the JCF are given in non-increasing order of their sizes. In 1951 Flanders published the following result [4, Theorem 2]:

Theorem 1.1. *If $A \in \mathbb{C}^{m \times n}$ and $B \in \mathbb{C}^{n \times m}$, then the JCFs of AB and BA may differ only in the sizes of the Jordan blocks at 0. Moreover, the difference between two corresponding sizes is at most one. Conversely, if the JCFs of $M \in \mathbb{C}^{m \times m}$ and $N \in \mathbb{C}^{n \times n}$ satisfy these properties, then $M = AB$ and $N = BA$, for some A, B .*

Theorem 1.1 has been revisited several times and re-proved using different techniques [1, 10, 11, 12, 14, 16]. In this paper, we investigate what happens if, instead of two matrices, we have products of k matrices, $A_1, A_2, \dots, A_k \in \mathbb{C}^{n \times n}$. We refer to products of A_1, \dots, A_k , in any order and with no repetitions of the factors, as *permuted products*.

We assume A_1, \dots, A_k are all $n \times n$ to ensure all permuted products are well defined. An important difference between $k = 2$ and $k > 2$ is that, without any assumption on A_1, \dots, A_k , the products of A_1, \dots, A_k have, in general, completely different eigenvalues for different permutations. One exception is the eigenvalue 0: if 0 is an eigenvalue of some product of A_1, \dots, A_k , then it must be an eigenvalue of every other product of A_1, \dots, A_k . Indeed, in Theorem 1.1 the eigenvalue 0 is treated exceptionally, with nontrivial results on the sizes of the Jordan blocks at 0. However, the following simple example with $k = 3$ shows that the difference between the sizes of Jordan blocks at $\lambda = 0$ can be arbitrarily large.

Example 1.2. *Let $A = \text{diag}(1, 1/2, \dots, 1/n)$, $B = -J_n(-1)^T$, and $C = (AB)^{-1}J_n(0)$, where $J_n(\lambda)$ is the $n \times n$ Jordan block at the eigenvalue λ [8, Def. 3.1.1]. Then $ABC = J_n(0)$ by construction, whereas CBA has a simple eigenvalue 0 and $n - 1$ nonzero eigenvalues.*

To verify the last statement, observe that B^{-1} is lower triangular with all elements on and below the main diagonal equal to 1. Therefore, the last two rows of $B^{-1}A^{-1}$ are equal, up to the last-but-one entry. Hence, the last two rows of $C = B^{-1}A^{-1}J_n(0)$ are equal. Since CBA is a product of $J_n(0)$ with some nonsingular matrices, we have $\text{rank } CBA = n - 1$. Moreover, the vector $v_0 = [1, 2, \dots, n]^T$ belongs to its kernel because $CBAv_0 = CB[1, 1, \dots, 1]^T = Ce_1 = 0$. Now suppose that there is a Jordan chain, so let v_1 be such that $CBAv_1 = v_0$. Then, since BA is invertible, there exists w such that $Cw = v_0$, but this is impossible as the elements of v_0 are all different while the last two rows of C are identical. So by contradiction 0 must be a simple eigenvalue of CBA .

Example 1.2 shows that it may be difficult to characterize the eigenvalues or Jordan block sizes for products of three or more matrices.

In [3], Fiedler introduced a decomposition of an $n \times n$ companion matrix into a product of $k = n$ matrices, $C = \prod_{i=1}^n M_i$, and showed that the product of the matrices M_i in any order is similar to C , hence all permuted products have the same JCF. For the nonzero eigenvalues, this is precisely what happens when $k = 2$, by Theorem 1.1. This motivates us to examine general conditions that allow an extension of Theorem 1.1 for nonzero eigenvalues to the case $k > 2$. The Fiedler factors M_i have the following properties:

(F1) Commutativity: $M_i M_j = M_j M_i$, if $|i - j| > 1$.

(F2) M_i are all nonsingular, possibly except for M_n .

Fiedler's results suggest the possibility of extending Theorem 1.1 to products of three or more matrices under appropriate commutativity conditions. Indeed, we will show that if the graph of non-commutativity relations is a forest (see Section 4), then all permuted products have the same Jordan blocks for nonzero eigenvalues. This commutativity assumption generalizes condition (F1), and imposes no requirement when $k = 2$, i.e., the two matrices can be arbitrary, thus recovering Theorem 1.1. We impose no nonsingularity condition such as (F2) because this imposes similarity, i.e., also the Jordan blocks at zero must be the same: an undesirable restriction given our goal of generalizing Flanders' theorem. Indeed, Theorem 1.1 shows that the difference in the sizes of Jordan blocks at zero is at most 1 when $k = 2$. One key result of our paper is that, for general k , under our commutativity conditions this difference is bounded by $k - 1$, and the bound is attainable.

For products of three matrices, our condition reduces to the requirement that one pair commutes, and we prove that the allowable sizes are exhaustive. More precisely, we prove that given two lists of these allowable sizes, there are matrices A, B, C such that the JCFs of ABC and CBA consist of Jordan blocks at $\lambda = 0$ whose sizes match those in the respective lists.

Several previous papers have addressed extensions of Flanders' result to many matrices. For example, [7] examines cyclic permutations and [5] derives conditions for the products to have the same trace, the same characteristic polynomial, or the same JCF, with focus on $k = 3$ or 2×2 matrices. Unlike in previous studies, we deal more thoroughly with *any* permutation and arbitrary n and $k \geq 3$, and work with commutativity conditions guaranteeing that the Jordan structures for nonzero eigenvalues coincide for all permutations.

The paper is organized as follows. Section 2 reviews basic notions and previous results. In Section 3 we analyze permuted products of $k = 3$ matrices. Section 4 discusses the case $k > 3$, which requires the use of permutations and graphs. We conclude in Section 5 with a summary and some open problems related to this work.

2. Notation, definitions and some consequences of Flanders' theorem

We follow the standard notation I_n and 0_n to denote, respectively, the $n \times n$ identity and null matrices. Given a square matrix $M \in \mathbb{C}^{n \times n}$, $\Lambda(M)$ denotes the spectrum (set of eigenvalues counting multiplicity) of M ; $\text{diag}(A_1, \dots, A_k)$ is the block-diagonal matrix whose diagonal blocks are A_1, \dots, A_k , in this order (that is, the *direct sum* of A_1, \dots, A_k). Two matrices $M, N \in \mathbb{C}^{n \times n}$ are *similar* if there is an invertible matrix P such that $PMP^{-1} = N$.

The *Jordan block of size $k \in \mathbb{N}$ at zero* is the $k \times k$ matrix

$$J_k(0) := \begin{bmatrix} 0 & 1 & & & \\ & \ddots & \ddots & & \\ & & 0 & 1 & \\ & & & & 0 \end{bmatrix},$$

and the *Jordan block of size k at $\lambda \in \mathbb{C}$* is the $k \times k$ matrix $J_k(\lambda) := J_k(0) + \lambda I_k$.

For a given $\lambda \in \mathbb{C}$, the *Segré characteristic of M at λ* , denoted by $\mathcal{S}_\lambda(M)$, is the list of the sizes of the Jordan blocks at λ in the JCF of M . In this paper we regard it as an infinite nonincreasing sequence of nonnegative integers, by attaching an infinite sequence of zeros at the end. Note that the Segré characteristic at any λ is uniquely determined, and that this definition includes also those complex numbers that are not eigenvalues of M , though in this case all entries in the Segré characteristic are zeros.

We use boldface for lists of nonnegative integers. Given two (possibly infinite) sequences of integers $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots)$ and $\boldsymbol{\mu}' = (\mu'_1, \mu'_2, \dots)$, we will often refer to the standard ℓ^∞ and ℓ^1 norms, which we denote by $\|\cdot\|_\infty$ and $\|\cdot\|_1$.

Given k matrices $A_1, \dots, A_k \in \mathbb{C}^{n \times n}$, by a *permuted product of A_1, \dots, A_k* we mean any of the products of A_1, \dots, A_k in all possible orders, without repetitions. The set of permuted products of A_1, \dots, A_k is denoted by $\mathcal{P}(A_1, \dots, A_k)$. For instance, for three matrices A, B, C , we have

$$\mathcal{P}(A, B, C) = \{ABC, ACB, BAC, BCA, CAB, CBA\}.$$

We will generally use the Π symbol to denote elements of $\mathcal{P}(A_1, \dots, A_k)$.

The following definition relates matrices M, N in Theorem 1.1, and plays a central role in this paper.

Definition 2.1. *A pair of matrices (M, N) , with $M \in \mathbb{C}^{m \times m}$ and $N \in \mathbb{C}^{n \times n}$, is a Flanders pair if there are two matrices $A \in \mathbb{C}^{m \times n}$ and $B \in \mathbb{C}^{n \times m}$ such that $M = AB$ and $N = BA$. In this case, we say that there is a Flanders bridge between M and N .*

We have the following elementary result:

Lemma 2.2. *If $M, N \in \mathbb{C}^{n \times n}$ are similar, then (M, N) is a Flanders pair.*

Proof. If $PMP^{-1} = N$, with P nonsingular, then we may take $B = PM$, $A = P^{-1}$, which satisfy $AB = M$ and $BA = N$. \square

The converse of Lemma 2.2 is not true in general. This is an immediate consequence of Theorem 1.1, since two matrices in a Flanders pair may have different Segré characteristic at zero and, as a consequence, different JCF. However, if M, N are nonsingular, then (M, N) is a Flanders pair if and only if M and N are similar. This is also an immediate consequence of Theorem 1.1.

The relation \mathcal{R} on $\mathbb{C}^{n \times n} \times \mathbb{C}^{n \times n}$ defined by “ MRN if (M, N) is a Flanders pair” is not an equivalence relation, since \mathcal{R} is not transitive. Moreover, Flanders pairs connecting three matrices M, N, Q in the form $(M, N), (N, Q)$ are closely related to our problem. The following direct consequence of Theorem 1.1 establishes some elementary features of these pairs.

Corollary 2.3. *If $M \in \mathbb{C}^{m \times m}$, $N \in \mathbb{C}^{n \times n}$ and $Q \in \mathbb{C}^{q \times q}$ are such that (M, N) and (N, Q) are Flanders pairs, then*

- (i) $\mathcal{S}_\lambda(M) = \mathcal{S}_\lambda(Q)$, for all $\lambda \neq 0$, and
- (ii) $\|\mathcal{S}_0(M) - \mathcal{S}_0(Q)\|_\infty \leq 2$.

In Corollary 3.6 we give a characterization of pairs of matrices M, Q as in the statement of Corollary 2.3 and with the same size. We will see, in particular, that, when M and Q have the same size, the converse of Corollary 2.3 also holds. Corollary 2.3 can be extended directly to more than three matrices.

Another feature of Theorem 1.1 we are interested in is its *exhaustivity*. The meaning of exhaustivity is exhibited in the following result.

Theorem 2.4. *Let $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots)$, and $\boldsymbol{\mu}' = (\mu'_1, \mu'_2, \dots)$ be two lists of integers with $\mu_1 \geq \mu_2 \geq \dots \geq 0$, and $\mu'_1 \geq \mu'_2 \geq \dots \geq 0$, such that*

- (i) $\|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_\infty \leq 1$, and
- (ii) $\|\boldsymbol{\mu}\|_1 = m$, $\|\boldsymbol{\mu}'\|_1 = n$.

Then, there exist two matrices $A \in \mathbb{C}^{m \times n}$ and $B \in \mathbb{C}^{n \times m}$, such that $\mathcal{S}_0(AB) = \boldsymbol{\mu}$ and $\mathcal{S}_0(BA) = \boldsymbol{\mu}'$.

Theorem 2.4 follows immediately from Theorem 1.1 just by noticing that it is always possible to construct two matrices $M \in \mathbb{C}^{m \times m}$ and $N \in \mathbb{C}^{n \times n}$ such that $\mathcal{S}_0(M) = \boldsymbol{\mu}$ and $\mathcal{S}_0(N) = \boldsymbol{\mu}'$, with $m, n, \boldsymbol{\mu}, \boldsymbol{\mu}'$ as in the statement of Theorem 2.4. It can be proved also in a direct way by explicitly constructing A and B . This is the approach followed in [12, Th. 3.3]. We present an extension of Theorem 2.4 to three matrices in Theorem 4.15. Our approach owes very much to the one in [12].

3. The case of three matrices

Unlike what happens for two matrices, given three matrices, $A, B, C \in \mathbb{C}^{n \times n}$, the spectra of two different permuted products of A, B, C may be completely different. To verify this, one may just take three random matrices A, B, C and compute the eigenvalues of ABC and ACB . This is related to the fact that

two similar matrices, as BC and CB are if one of B, C is nonsingular, may give two matrices with completely different spectra when multiplied on the left by a third matrix A . To what extent may the spectra of different permutation products of three given matrices differ? One restriction is that the determinants must all be the same, which implies that if 0 is an eigenvalue of some permuted product then it must be shared by all permuted products. However, as we have seen in Example 1.2, the Jordan structure of the eigenvalue 0 may differ from one product to another. Let us first consider the case of nonsingular matrices. The following result shows that, without any additional assumptions, the only restriction on the spectra of permuted products of three nonsingular matrices A, B, C is that they all have the same determinant. It is a restatement, with a more straightforward proof, of Theorem 4 in [6].

Theorem 3.1. *Let $\Lambda_1 = \{\lambda_{11}, \dots, \lambda_{n1}\}$ and $\Lambda_2 = \{\lambda_{12}, \dots, \lambda_{n2}\}$ be two sets of n nonzero complex numbers, with possible repetitions. If $\lambda_{11} \cdots \lambda_{n1} = \lambda_{12} \cdots \lambda_{n2}$, then there are three matrices $A, B, C \in \mathbb{C}^{n \times n}$, such that $\Lambda(ABC) = \Lambda_1$ and $\Lambda(ACB) = \Lambda_2$.*

Proof. By Lemma 2.2, it suffices to find two similar matrices $M, N \in \mathbb{C}^{n \times n}$, and a third matrix $A \in \mathbb{C}^{n \times n}$, such that $\Lambda(AM) = \Lambda_1$ and $\Lambda(AN) = \Lambda_2$. This can be done using only diagonal matrices. More precisely, set $r_1 \neq 0$ (arbitrary), $a_1 := \lambda_{11}/r_1$ and, recursively for $i = 2, \dots, n$, $r_i := \lambda_{i2}/a_{i-1}$, $a_i := \lambda_{i1}/r_i$. Note that, with these definitions, we have

$$a_n r_1 = \frac{(a_1 r_1)(a_2 r_2) \cdots (a_n r_n)}{(a_1 r_2)(a_2 r_3) \cdots (a_{n-1} r_n)} = \frac{\lambda_{11} \lambda_{21} \cdots \lambda_{n1}}{\lambda_{12} \lambda_{22} \cdots \lambda_{n-1,2}} = \lambda_{n2}.$$

Hence, if we set $M = \text{diag}(r_1, r_2, \dots, r_n)$, $N = \text{diag}(r_2, r_3, \dots, r_n, r_1)$, and $A = \text{diag}(a_1, \dots, a_n)$, then M is similar to N , and $AM = \text{diag}(\lambda_{11}, \dots, \lambda_{n1})$, $AN = \text{diag}(\lambda_{12}, \dots, \lambda_{n2})$, as required. \square

Under the conditions of the statement of Theorem 3.1, by Theorem 1.1 we have $\Lambda(ABC) = \Lambda(CAB) = \Lambda(BCA) = \Lambda_1$, and $\Lambda(ACB) = \Lambda(BAC) = \Lambda(CBA) = \Lambda_2$. Moreover, as a consequence of Theorem 1.1, the set of permuted products is partitioned into two classes, namely: $\mathcal{C}_1 = \{ABC, BCA, CAB\}$, and $\mathcal{C}_2 = \{ACB, BAC, CBA\}$. Any two products in each class are related by a ‘‘cyclic permutation’’, so they form a Flanders pair. Hence, we can relate the JCFs of these permuted products. The remaining question is to relate the JCFs of permuted products in \mathcal{C}_1 with the ones in \mathcal{C}_2 . Theorem 3.1 shows that, if A, B, C are nonsingular, there may be no relationship at all between the spectra of products in different classes.

Motivated by the work of Fiedler, here we require that at least two of A, B, C commute. As we see in Section 4, if we consider formal products of an arbitrary number of matrices, commutativity conditions allow us to characterize those cases where any two arbitrary permutations are linked by a sequence of Flanders bridges. In this case, all permuted products have the same Segré characteristic at each nonzero complex number.

Proposition 3.2. *Let $A, B, C \in \mathbb{C}^{n \times n}$ be such that at least two of A, B, C commute. Let $\Pi_1, \Pi_2 \in \mathcal{P}(A, B, C)$. Then*

- (i) $\mathcal{S}_\lambda(\Pi_1) = \mathcal{S}_\lambda(\Pi_2)$, for all $\lambda \neq 0$, and
- (ii) $\|\mathcal{S}_0(\Pi_1), \mathcal{S}_0(\Pi_2)\|_\infty \leq 2$.

Proof. By Corollary 2.3, it suffices to show that, in the conditions of the statement, one of the following situations occurs:

1. (Π_1, Π_2) is a Flanders pair.
2. There exists $\tilde{\Pi} \in \mathcal{P}(A, B, C)$ such that $(\Pi_1, \tilde{\Pi})$ and $(\tilde{\Pi}, \Pi_2)$ are Flanders pairs.

In the conditions of the statement there are at most 4 distinct elements in $\mathcal{P}(A, B, C)$, which give at most 6 distinct (non-ordered) pairs of permuted products. Let us assume, without loss of generality, that $AC = CA$. In this case, the elements in $\mathcal{P}(A, B, C)$ (including Π_1 and Π_2) are ABC, ACB, BAC, CBA , and $(ABC, ACB), (ABC, BAC), (ACB, BAC)$ and (BAC, CBA) are Flanders pairs. Hence, one of the situations described above holds for Π_1 and Π_2 . \square

The following technical Lemma 3.3 is used to prove Theorem 4.15:

Lemma 3.3. *Let $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots)$, $\boldsymbol{\mu}' = (\mu'_1, \mu'_2, \dots) \in \ell^1$ be two sequences of nonnegative integers. Suppose that*

- (i) $\|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_\infty = 2$, and
- (ii) $\|\boldsymbol{\mu}\|_1 = \|\boldsymbol{\mu}'\|_1 = n$.

Then we may rearrange $\boldsymbol{\mu}$ and $\boldsymbol{\mu}'$ in such a way that

$$\boldsymbol{\mu} = (\mu_{i_1}, \mu_{i_2}, \mu_{i_3}; \mu_{i_4}, \mu_{i_5}, \mu_{i_6}; \dots), \quad \boldsymbol{\mu}' = (\mu'_{i_1}, \mu'_{i_2}, \mu'_{i_3}; \mu'_{i_4}, \mu'_{i_5}, \mu'_{i_6}; \dots),$$

with

$$\mu_{i_j} + \mu_{i_{j+1}} + \mu_{i_{j+2}} = \mu'_{i_j} + \mu'_{i_{j+1}} + \mu'_{i_{j+2}}, \quad \text{for all } j \equiv 1 \pmod{3}. \quad (3.1)$$

Proof. Let $m = \max(\|\boldsymbol{\mu}\|_0, \|\boldsymbol{\mu}'\|_0)$ be the maximum of the number of nonzero elements in $\boldsymbol{\mu}$ and $\boldsymbol{\mu}'$. We may assume that they both have the same length, by adding zeros to one of them if necessary. The proof is carried out by induction on m . For $m \leq 3$ the result is trivial. Suppose the result holds for lengths up to $m - 1$, and let us prove it for length equal to m . By condition (i) in the statement, there is some $i \geq 0$ such that $|\mu_i - \mu'_i| = 2$. Without loss of generality we may assume that $\mu_i = \mu'_i + 2$. Now, condition (ii) in the statement implies that at least one of the following situations must occur:

- (A1) There is some $j \geq 0$ such that $\mu'_j = \mu_j + 2$, or
- (A2) There are some $k, \ell \geq 0$, with $k \neq \ell$, such that $\mu'_k = \mu_k + 1$ and $\mu'_\ell = \mu_\ell + 1$.

In case (A1), we may rearrange $\boldsymbol{\mu}$ and $\boldsymbol{\mu}'$, by adding one extra zero in each list, in the form:

$$\begin{aligned}\boldsymbol{\mu} &= (\mu_i, \mu_j, 0; \tilde{\boldsymbol{\mu}}), \\ \boldsymbol{\mu}' &= (\mu'_i, \mu'_j, 0; \tilde{\boldsymbol{\mu}}'),\end{aligned}$$

where $\tilde{\boldsymbol{\mu}}$ and $\tilde{\boldsymbol{\mu}}'$ are obtained from $\boldsymbol{\mu}$ and $\boldsymbol{\mu}'$, respectively, by removing the i th and j th elements. Now, the result follows by the induction hypothesis on $\tilde{\boldsymbol{\mu}}$ and $\tilde{\boldsymbol{\mu}}'$.

In case (A2) we may rearrange:

$$\begin{aligned}\boldsymbol{\mu} &= (\mu_i, \mu_k, \mu_\ell; \tilde{\boldsymbol{\mu}}), \\ \boldsymbol{\mu}' &= (\mu'_i, \mu'_k, \mu'_\ell; \tilde{\boldsymbol{\mu}}'),\end{aligned}$$

where $\tilde{\boldsymbol{\mu}}$ and $\tilde{\boldsymbol{\mu}}'$ are obtained from $\boldsymbol{\mu}$ and $\boldsymbol{\mu}'$, respectively, by removing the i th, k th and ℓ th elements. Again, the result follows by induction on $\tilde{\boldsymbol{\mu}}$ and $\tilde{\boldsymbol{\mu}}'$. \square

The main result of this section is an extension of [12, Th. 3.3] to three matrices A, B, C under the commutativity condition $AC = CA$.

Theorem 3.4. *Let $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots, 0, \dots)$, $\boldsymbol{\mu}' = (\mu'_1, \mu'_2, \dots, 0, \dots) \in \ell^1$ be two nonincreasing sequences of nonnegative integers such that*

$$(i) \quad \|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_\infty \leq 2, \text{ and}$$

$$(ii) \quad \|\boldsymbol{\mu}\|_1 = \|\boldsymbol{\mu}'\|_1 = n.$$

Then, there exist three matrices $A, B, C \in \mathbb{C}^{n \times n}$, such that $AC = CA$ and

$$\mathcal{S}_0(ABC) = \boldsymbol{\mu}, \quad \text{and} \quad \mathcal{S}_0(CBA) = \boldsymbol{\mu}'.$$

Proof. First, notice that if $\|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_\infty \leq 1$, then by Theorem 1.1 there exist $A, B \in \mathbb{C}^{n \times n}$ such that $\mathcal{S}_0(AB) = \boldsymbol{\mu}$ and $\mathcal{S}_0(BA) = \boldsymbol{\mu}'$. In this case, we may take $C = I_n$ and we are done. Hence it remains to consider the case $\|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_\infty = 2$. The proof reduces to showing that the statement is true in the following two cases:

$$(A1) \quad \boldsymbol{\mu} = (m, n, 0, \dots), \quad \boldsymbol{\mu}' = (m - 2, n + 2, 0, \dots)$$

$$(A2) \quad \boldsymbol{\mu} = (m, n, q, 0, \dots), \quad \boldsymbol{\mu}' = (m - 2, n + 1, q + 1, 0, \dots),$$

with $m, n, q \geq 0$ and $m \geq 2$. Indeed, let us assume that the result is true for both (A1) and (A2), and let $\boldsymbol{\mu}$ and $\boldsymbol{\mu}'$ be as in the statement. By Lemma 3.3, we can rearrange $\boldsymbol{\mu}$ and $\boldsymbol{\mu}'$ in such a way that they are partitioned as

$$\boldsymbol{\mu} = (\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_\alpha, 0, \dots), \quad \text{and} \quad \boldsymbol{\mu}' = (\boldsymbol{\mu}'_1, \dots, \boldsymbol{\mu}'_\alpha, 0, \dots),$$

where the pairs $(\boldsymbol{\mu}_i, \boldsymbol{\mu}'_i)$ for $i = 1, \dots, \alpha$ are such that $\|\boldsymbol{\mu}_i\|_1 = \|\boldsymbol{\mu}'_i\|_1 =: n_i$ and they either satisfy $\|\boldsymbol{\mu}_i - \boldsymbol{\mu}'_i\|_\infty \leq 1$ or are of one of the forms (A1), (A2). Now, since the result is true for both (A1) and (A2), and also for tuples of distance at most 1, there are matrices $A_1, B_1, C_1 \in \mathbb{C}^{n_1 \times n_1}, \dots, A_\alpha, B_\alpha, C_\alpha \in \mathbb{C}^{n_\alpha \times n_\alpha}$,

such that $A_i C_i = C_i A_i$, and $\mathcal{S}_0(A_i B_i C_i) = (\boldsymbol{\mu}_i, 0, \dots)$, $\mathcal{S}_0(C_i B_i A_i) = (\boldsymbol{\mu}'_i, 0, \dots)$, for $i = 1, \dots, \alpha$. Then the matrices

$$A = \text{diag}(A_1, \dots, A_\alpha), \quad B = \text{diag}(B_1, \dots, B_\alpha), \quad C = \text{diag}(C_1, \dots, C_\alpha)$$

satisfy $AC = CA$ and $\mathcal{S}_0(ABC) = \boldsymbol{\mu}$, $\mathcal{S}_0(CBA) = \boldsymbol{\mu}'$.

It remains to prove that the result is true in cases (A1) and (A2). Consider (A1) first. Denote by E_{ij} the matrix, of the appropriate size, whose (i, j) entry is equal to 1 and the remaining entries are zero. Set

$$A = \text{diag}(I_{m-1}, 0, I_n), \quad B = J_{m+n}(0) + E_{m+n,1}, \quad C = \text{diag}(0, I_{m+n-1}).$$

Clearly we have $AC = CA$. Direct computation gives $ABC = \text{diag}(J_m(0), J_n(0))$, and $CBA = \text{diag}(0, J_{m-2}(0), J_{n+1}(0)) + E_{m+n,1}$. Now, $\text{diag}(0, J_{n+1}(0)) + E_{n+2,1}$ is similar to $J_{n+2}(0)$, because its only eigenvalue is 0 and its rank deficiency is one. Consequently, the JCF of CBA is $\text{diag}(J_{m-2}(0), J_{n+2}(0))$, so $\mathcal{S}_0(ABC) = (m, n, 0, \dots)$ and $\mathcal{S}_0(CBA) = (m-2, n+2, 0, \dots)$, as required.

Next consider (A2). Set

$$A = \text{diag}(0, I_{m+n+q-1}), \quad C = \text{diag}(I_{n+q+1}, 0, I_{m-2}),$$

for which $AC = CA$, and set also

$$B = \text{diag}(J_{q+1}(0), J_{m+n-1}(0)) + E_{m+n+q,1}.$$

Direct computation gives

$$ABC = \text{diag}(0, J_q(0), J_n(0), J_{m-1}(0)) + E_{m+n+q,1},$$

and

$$CBA = \text{diag}(J_{q+1}(0), J_{n+1}(0), J_{m-2}(0)).$$

Note that $\text{diag}(0, J_q(0), J_n(0), J_{m+1}(0)) + E_{m+n+q,1}$ is permutation similar to $\text{diag}(J_q(0), J_n(0), \text{diag}(0, J_{m-1}(0)) + E_{m,1})$. Since, as before, $\text{diag}(0, J_{m-1}(0)) + E_{m,1}$ is similar to $J_m(0)$, we conclude that $\mathcal{S}_0(ABC) = (m, n, q, 0, \dots)$ and $\mathcal{S}_0(CBA) = (m-2, n+1, q+1, 0, \dots)$, as required. \square

Remark 3.5. If $\|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_\infty = 2$, then the matrices A, B, C constructed in the proof of Theorem 4.15 have the property that neither of the pairs (A, B) and (B, C) commutes, so there is exactly one commutativity relation in this case. In graph theoretical terminology (see Section 4), the graph of non-commutativity relations is a tree.

Our last result in this section concerns the “non-transitivity” of Flanders pairs.

Corollary 3.6. *Let $M, Q \in \mathbb{C}^{n \times n}$. Then, the following conditions are equivalent:*

- (a) *There exists $N \in \mathbb{C}^{n \times n}$ such that (M, N) and (N, Q) are Flanders pairs.*

- (b) $\mathcal{S}_\lambda(M) = \mathcal{S}_\lambda(Q)$, for all $\lambda \neq 0$, and $\|\mathcal{S}_0(M) - \mathcal{S}_0(Q)\|_\infty \leq 2$.
- (c) *There are three matrices $A, B, C \in \mathbb{C}^{n \times n}$ such that $AC = CA$, M is similar to ABC , and Q is similar to CBA .*

Proof. The implication (a) \Rightarrow (b) is Corollary 2.3. Suppose that (b) holds. Without loss of generality, we may assume that M and Q are given in JCF, so that $M = \text{diag}(M_r, M_s)$, and $Q = \text{diag}(Q_r, Q_s)$ where M_r, Q_r contain Jordan blocks associated with nonzero eigenvalues, and M_s, Q_s are Jordan blocks for $\lambda = 0$. By hypothesis, we have $M_r = Q_r$ and $\|\mathcal{S}_0(M_s) - \mathcal{S}_0(Q_s)\|_\infty \leq 2$. Using Theorem 4.15 with $\boldsymbol{\mu} = \mathcal{S}_0(M_s)$ and $\boldsymbol{\mu}' = \mathcal{S}_0(Q_s)$, we see that there exist A_s, B_s, C_s such that $A_s C_s = C_s A_s$, $A_s B_s C_s = M_s$, and $C_s B_s A_s = Q_s$. The block diagonal matrices $A = \text{diag}(I_m, A_s)$, $B = \text{diag}(M_r, B_s)$, $C = \text{diag}(I_m, C_s)$, where m is the size of both M_r and Q_r , fulfill the conditions in (c).

Finally, suppose that (c) holds. Let $N = BCA$. Then (M, N) is clearly a Flanders pair and, by the condition $AC = CA$, so is the pair (N, Q) . \square

We want to emphasize the difference between Corollary 3.6 and Theorem 1.1. The natural extension of Theorem 1.1 to three matrices would be that (M, N) and (N, Q) are Flanders pairs if and only if there are three matrices A, B, C such that $AC = CA$ and $M = ABC$, $Q = CBA$. However, we have not been able to prove that this is true and we have not found a counterexample. This issue remains an open problem (see Open Problem 3 in Section 5).

4. More than three matrices

For permutations in Σ_k , the symmetric group of $\{1, \dots, k\}$, we use the cyclic notation $\sigma = (i_1 i_2 \dots i_s)$ to mean that $\sigma(i_j) = i_{j+1}$, for $j = 1, \dots, s-1$, $\sigma(i_s) = i_1$, and $\sigma(i) = i$, for $i \neq i_1, \dots, i_s$.

An element in $\mathcal{P}(A_1, \dots, A_k)$ is related to a permutation $\sigma \in \Sigma_k$ in the form $A_{\sigma^{-1}(1)} A_{\sigma^{-1}(2)} \dots A_{\sigma^{-1}(k)}$, that is, $\sigma(i)$ is the position of the factor A_i in the permuted product. In this case, we write $\Pi_\sigma := A_{\sigma^{-1}(1)} A_{\sigma^{-1}(2)} \dots A_{\sigma^{-1}(k)}$.

Definition 4.1. *Given a permutation $\sigma \in \Sigma_k$, a cyclic permutation of σ is a permutation of the form $(1 \ 2 \ \dots \ k)^\ell \sigma$, for some $\ell \geq 0$. We say that σ, τ are cyclically related if τ is a cyclic permutation of σ .*

Accordingly, given a permuted product $\Pi_\sigma \in \mathcal{P}(A_1, \dots, A_k)$, a cyclic permutation of Π_σ is a permuted product of the form $\Pi_\tau \in \mathcal{P}(A_1, \dots, A_k)$, with $\tau = (1 \ 2 \ \dots \ k)^\ell \sigma$, for some $\ell \in \mathbb{N}$. If Π_σ is a cyclic permutation of Π_τ , then Π_σ and Π_τ are cyclically equivalent, and we write $\Pi_\sigma \sim_C \Pi_\tau$.

We note that \sim_C is, indeed, an equivalence relation. Moreover, if $\Pi_{\sigma_1} \sim_C \Pi_{\sigma_2}$, then $(\Pi_{\sigma_1}, \Pi_{\sigma_2})$ is a Flanders pair. Conversely, if $(\Pi_{\sigma_1}, \Pi_{\sigma_2})$ is a Flanders pair for all A_1, \dots, A_k (that is, as a “formal product”), then Π_{σ_1} is a cyclic permutation of Π_{σ_2} .

Definition 4.2. Given two permutations $\sigma_1, \sigma_2 \in \Sigma_k$, we say that i_1, \dots, i_g , with $1 \leq i_1, \dots, i_g \leq k$, have the same order in σ_1 and σ_2 up to cyclic permutations if i_1, \dots, i_g appear in the same order in $\tilde{\sigma}_1 := (1 \ 2 \ \dots \ k)^\alpha \sigma_1$ and σ_2 for some $\alpha \geq 0$.

Accordingly, given $\Pi_{\sigma_1}, \Pi_{\sigma_2} \in \mathcal{P}(A_1, \dots, A_k)$, we say that A_{i_1}, \dots, A_{i_g} have the same cyclic order in both Π_{σ_1} and Π_{σ_2} if i_1, \dots, i_g have the same order in σ_1 and σ_2 up to cyclic permutations.

4.1. Inverse eigenvalue problem

We start with an observation that characterizes Σ_k up to cyclic permutations.

Lemma 4.3. Let $\sigma, \tau \in \Sigma_k$ be two permutations. Then σ and τ are cyclically related if and only if all triples i_1, i_2, i_3 , with $1 \leq i_1, i_2, i_3 \leq k$ have the same order in σ and τ up to cyclic permutations.

Proof. If $\sigma = (1 \ 2 \ \dots \ k)^\ell \tau$, for some $\ell \geq 0$, then it is clear that each triple i_1, i_2, i_3 has the same order up to cyclic permutations in both σ and τ .

Conversely, assume that every triple i_1, i_2, i_3 has the same order up to cyclic permutations in σ and τ . Let $\alpha, \beta \geq 0$ be such that $\tilde{\sigma} := (1 \ 2 \ \dots \ k)^\alpha \sigma$ and $\tilde{\tau} := (1 \ 2 \ \dots \ k)^\beta \tau$ satisfy $\tilde{\sigma}(1) = 1 = \tilde{\tau}(1)$. Suppose $\tilde{\sigma} \neq \tilde{\tau}$ and let $\nu = \min\{i : \tilde{\sigma}(i) \neq \tilde{\tau}(i)\}$. Then $1, \tilde{\sigma}(\nu), \tilde{\tau}(\nu)$ do not have the same order up to cyclic permutations in $\tilde{\sigma}$ and $\tilde{\tau}$, a contradiction. Hence, σ and τ are cyclically related. \square

We next show that it is possible that *any* two permuted products Π_1, Π_2 have different spectra unless $\Pi_1 \sim_C \Pi_2$.

Proposition 4.4. For each $k \geq 3$, there exist matrices, $A_1, \dots, A_k \in \mathbb{C}^{n \times n}$ such that for any two permuted products Π_1 and Π_2 belonging to different equivalence classes of $\mathcal{P}(A_1, \dots, A_k)$ under \sim_C , $\Lambda(\Pi_1)$ and $\Lambda(\Pi_2)$ are different.

Proof. First, let us order all the $\binom{k}{3}$ triples (i_1, i_2, i_3) , with $1 \leq i_1 < i_2 < i_3 \leq k$ using, for instance, the lexicographic order. This order induces an ordered list of length $3 \cdot \binom{k}{3} = \frac{k(k-1)(k-2)}{2}$, denoted by \mathcal{L} , after adjoining all triples in the given order. For instance, for $k = 4$ we get the list: $\mathcal{L} = (1, 2, 3; 1, 2, 4; 1, 3, 4; 2, 3, 4)$. Now, let $\gamma : \left\{1, 2, \dots, \frac{k(k-1)(k-2)}{2}\right\} \rightarrow \{1, 2, \dots, k\}$ be the map defined by $\gamma(i) = \mathcal{L}_i$ (the i th number in \mathcal{L}). For each $j = 1, \dots, \binom{k}{3}$, by Theorem 3.1, there are three matrices $\tilde{A}_{3j-2}, \tilde{A}_{3j-1}, \tilde{A}_{3j} \in \mathbb{C}^{2 \times 2}$, such that $\Lambda(\tilde{A}_{3j-2} \tilde{A}_{3j-1} \tilde{A}_{3j}) \neq \Lambda(\tilde{A}_{3j} \tilde{A}_{3j-1} \tilde{A}_{3j-2})$. For $i = 1, \dots, k$, define

$$A_i = \text{diag} \left(A_{i1}, A_{i2}, \dots, A_{i, \binom{k}{3}} \right) \in \mathbb{C}^{2 \binom{k}{3} \times 2 \binom{k}{3}},$$

where

$$A_{ij} = \begin{cases} \tilde{A}_{3(j-1)+r}, & \text{if there is some } 1 \leq r \leq 3 \text{ such that } \gamma(3(j-1)+r) = i, \\ I_2, & \text{otherwise.} \end{cases}$$

For instance, for $k = 4$ we have $A_1 = \text{diag}(\tilde{A}_1, \tilde{A}_4, \tilde{A}_7, I_2)$, $A_2 = \text{diag}(\tilde{A}_2, \tilde{A}_5, I_2, \tilde{A}_{10})$, $A_3 = \text{diag}(\tilde{A}_3, I_2, \tilde{A}_8, \tilde{A}_{11})$, $A_4 = \text{diag}(I_2, \tilde{A}_6, \tilde{A}_9, \tilde{A}_{12})$.

Let Π_{σ_1} and Π_{σ_2} be two permuted products in $\mathcal{P}(A_1, \dots, A_k)$ that are not cyclically equivalent. By Lemma 4.3, there is a triple (i_1, i_2, i_3) , with $1 \leq i_1, i_2, i_3 \leq k$, such that i_1, i_2, i_3 appear in this order in σ_1 , and they appear in the order i_3, i_2, i_1 in σ_2 , up to cyclic permutations. The triple (i_1, i_2, i_3) corresponds to a triple $(3j-2, 3j-1, 3j)$ in \mathcal{L} for some $j = 1, \dots, \binom{k}{3}$, such that $\Lambda(\tilde{A}_{3j-2}\tilde{A}_{3j-1}\tilde{A}_{3j}) \neq \Lambda(\tilde{A}_{3j}\tilde{A}_{3j-1}\tilde{A}_{3j-2})$. The result follows from the inclusions $\Lambda(\tilde{A}_{3j-2}\tilde{A}_{3j-1}\tilde{A}_{3j}) \subseteq \Lambda(\Pi_{\sigma_1})$ and $\Lambda(\tilde{A}_{3j}\tilde{A}_{3j-1}\tilde{A}_{3j-2}) \subseteq \Lambda(\Pi_{\sigma_2})$. \square

It is worth noting that, in the construction of the proof of Proposition 4.4, the spectra of Π_{σ_1} and Π_{σ_2} are not necessarily disjoint. Note also that the size of the matrices, namely $n = k(k-1)(k-2)/2$, depends on k .

All permuted products in $\mathcal{P}(A_1, \dots, A_k)$ have the same determinant. Equivalently, the product of their eigenvalues is the same for all permuted products. One may wonder whether this is the only restriction on the eigenvalues of permuted products belonging to different classes under cyclic permutations, as it is for three matrices. More generally, we may pose the following problem. Here and hereafter, for a given set Λ of complex numbers, the notation $\prod_{\lambda \in \Lambda} \lambda$ denotes the product of all numbers in Λ .

Inverse eigenvalue problem for permuted products of k matrices: *Given $(k-1)!$ sets of n nonzero complex numbers, $\Lambda_1, \dots, \Lambda_{(k-1)!}$, such that $\prod_{\lambda \in \Lambda_i} \lambda = \prod_{\lambda \in \Lambda_j} \lambda$, for all $1 \leq i, j \leq (k-1)!$, find matrices A_1, \dots, A_k , with $A_i \in \mathbb{C}^{n \times n}$, for $i = 1, \dots, k$, such that $\Lambda(\Pi_j) = \Lambda_j$, for $j = 1, \dots, (k-1)!$, where $\Pi_j \in \mathcal{P}(A_1, \dots, A_k)$ belongs to the j th equivalence class under \sim_C .*

In Section 3 we saw that the ‘‘Inverse eigenvalue problem for permuted products of $k = 3$ matrices’’ is always solvable. The following result states that this is not true for k large enough.

Theorem 4.5. *Let n, k be two integers such that $(k-1)!(n-1) + 1 > kn^2$. Then, there exist $(k-1)!$ sets of nonzero complex numbers $\Lambda_1, \dots, \Lambda_{(k-1)!}$, with $|\Lambda_i| = n$ and $\prod_{\lambda \in \Lambda_i} \lambda = \prod_{\lambda \in \Lambda_j} \lambda$, for all $1 \leq i, j \leq (k-1)!$, such that there are no matrices $A_1, \dots, A_k \in \mathbb{C}^{n \times n}$ satisfying $\Lambda(\Pi_j) = \Lambda_j$, for $j = 1, \dots, (k-1)!$, where $\Pi_j \in \mathcal{P}(A_1, \dots, A_k)$ belongs to the j th equivalence class under \sim_C .*

Proof. We first note that prescribing the eigenvalues of a matrix A is equivalent to prescribing the coefficients of the characteristic polynomial $p_A(\lambda) := \det(\lambda I - A)$. Let $A_1, \dots, A_k \in \mathbb{C}^{n \times n}$ be arbitrary matrices and let $X = \text{vec}([A_1, \dots, A_k]) \in \mathbb{C}^{kn^2 \times 1}$ be the *vectorization* of the block matrix $[A_1 \dots A_k]$ [9, Def. 4.2.9]. Let us denote by $\Pi_1, \dots, \Pi_{(k-1)!}$ the representatives of each of the equivalence classes of $\mathcal{P}(A_1, \dots, A_k)$ under \sim_C . Define the map

$$\begin{aligned} P: \mathbb{C}^M &\longrightarrow \mathbb{C}^N \\ X &\longmapsto P(X) = (P_1(X), \dots, P_N(X)), \end{aligned}$$

where $P(X)$ is the vector containing the coefficients of the characteristic polynomials of $\Pi_1, \dots, \Pi_{(k-1)!}$, in a certain pre-fixed order. P is a polynomial map, since the coefficients of the characteristic polynomial of a matrix are polynomial functions of the entries of the matrix. Moreover, we have $M = kn^2$ and $N = (k-1)!(n-1)+1$. Indeed, the necessary condition $\prod_{\lambda \in \Lambda_i} \lambda = \prod_{\lambda \in \Lambda_j} \lambda$, for $1 \leq i, j \leq (k-1)!$, is equivalent to the fact that the zero-degree coefficient of all characteristic polynomials of Π_j , for $j = 1, \dots, (k-1)!$, coincide. We may just slightly modify the definition of P , in such a way that, instead of n coefficients for each characteristic polynomial, we just have $(n-1)$ coefficients. Together with the choice of the determinant, this gives the $(k-1)!(n-1)+1$ coordinates in $P(X)$.

Now, the ‘‘Inverse eigenvalue problem for permuted products of k matrices with size $n \times n$ ’’ is solvable, for k and n , if and only if P is surjective for these k and n . It is known that a polynomial map from \mathbb{C}^M to \mathbb{C}^N is not surjective when $N > M$ [15, Th. 7, Ch. I, §6], so the result follows. \square

4.2. Graph theoretical description of $\mathcal{P}(A_1, \dots, A_k)$

We will see in Section 4.3 (see Theorem 4.16) how to characterize the maximum distance between the Segré characteristics of the zero eigenvalue in any two given products in $\mathcal{P}(A_1, \dots, A_n)$ using the graph of non-commutativity relations of pairs of matrices in $\{A_1, \dots, A_k\}$. From this graph, there arises an interesting combinatorial theory connected to this problem. The main result in this section is Theorem 4.13, which allows us to derive the main part of Theorem 4.16 as a direct consequence.

For the basic notions in graph theory we follow [2]. A *graph* is a pair of sets $\mathcal{G} = (V, E)$, where $V = \{1, \dots, k\}$ is the set of *vertices*, and E is the set of *edges*; an edge is a two element subset of V . Here $\{i, j\} \in E$ means that there is an edge *joining* i with j . By this definition, multiple edges between the same pair of vertices and edges joining a vertex to itself are disallowed.

A sequence of edges $\{i_0, i_1\}, \{i_1, i_2\}, \dots, \{i_{m-1}, i_m\}$ is called a *path* of length m if all vertices i_j are distinct. The sequence is called a *cycle of length* m if $m \geq 3$, $i_m = i_0$, and all vertices i_j , with $0 < j < m$, are different from each other and i_0 . We say that a graph has a cycle if a subset of its edges is a cycle. A graph $\mathcal{G} = (V, E)$ is *connected* if, for any pair of vertices in V , there is at least one path containing them. A *forest* is a graph that has no cycles, and a *tree* is a connected forest. The *degree* of a vertex $i \in V$ in the graph $\mathcal{G} = (V, E)$ is the number of vertices joined to i . A *leaf* is a vertex of degree one, and the *parent* of a leaf is the only vertex joined to it. A *cut* of $\mathcal{G} = (V, E)$ is a partition of $V = V_1 \cup V_2$ ($V_1 \cap V_2 = \emptyset$). Given a cut V_1, V_2 , we say that an edge $\{i, j\} \in E$ *crosses the cut* if i and j each lie in different V_b , $b = 1, 2$.

An *oriented graph* is a pair $\mathcal{G} = (V, E)$ where V is again a set of vertices and E is a set of ordered pairs of elements of V , that is, $E \subset V \times V$. Here $(i, j) \in E$ means that there is an edge *joining* i with j *from* i *to* j . A path (or cycle) of length m in a oriented graph is likewise a sequence of m elements of E of the form $(i_0, i_1), (i_1, i_2), \dots, (i_{m-1}, i_m)$, with distinct $\{i_j\}$ (except $i_m = i_0$ for cycles). An *acyclic* oriented graph is an oriented graph with no cycles.

An *orientation* of $\mathcal{G} = (V, E)$ is a function $\omega : E \rightarrow V \times V$ that assigns to each vertex $\{i, j\} \in E$ one of the ordered pairs (i, j) or (j, i) . Note that the set-wise image of E under ω , denoted $\omega(E) = \{\omega(e) : e \in E\}$, associates a graph \mathcal{G} with an oriented graph $(V, \omega(E))$. The orientation ω is said to be *acyclic* if $(V, \omega(E))$ has no cycles. The set of acyclic orientations of \mathcal{G} is denoted by $\mathcal{A}(\mathcal{G})$. Any total order \preceq of V determines an acyclic orientation $\omega \in \mathcal{A}(\mathcal{G})$ by $\omega(\{i, j\}) = (i, j)$ if and only if $i \prec j$ and $\{i, j\} \in E$. The converse is also true, as the following result shows.

Proposition 4.6. *Let $\mathcal{G} = (V, E)$ be a graph, and $\omega \in \mathcal{A}(\mathcal{G})$. Then, there is a total order \preceq of V such that ω is the acyclic orientation of \mathcal{G} determined by \preceq (a topological sort of ω).*

For a proof of Proposition 4.6 we refer the reader to [13, p. 137].

To motivate the following definition, the graph that is our primary concern is $\mathcal{G} = (V, E)$, where E encodes the non-commutativity relations on k matrices (see Definition 4.14). An edge $\{i, j\} \in E$ represents the fact that matrices A_i and A_j do not commute. Meanwhile, we continue to associate elements of Σ_k with elements of $\mathcal{P}(A_1, \dots, A_k)$ via Π_σ .

Definition 4.7. *Given a graph $\mathcal{G} = (V, E)$ with $V = \{1, 2, \dots, k\}$, we say that $\tau \circ \sigma$ is an allowed swap of $\sigma \in \Sigma_k$ when $\tau = (i \ i+1)$ is a transposition with $\{\sigma^{-1}(i), \sigma^{-1}(i+1)\} \notin E$, for some $1 \leq i \leq k-1$.*

The proof of the following result is straightforward.

Proposition 4.8. *Let $\mathcal{G} = (V, E)$ be a graph with $V = \{1, 2, \dots, k\}$. Let $\sim_{\mathcal{G}}$ be the relation on Σ_k defined by:*

$$\sigma_1 \sim_{\mathcal{G}} \sigma_2 \Leftrightarrow \sigma_2 = \tau_s \circ \dots \circ \tau_2 \circ \tau_1 \circ \sigma_1,$$

where, for each $i = 1, \dots, s$, τ_i is an allowed swap of $\tau_{i-1} \circ \dots \circ \tau_1 \circ \sigma_1$. Then $\sim_{\mathcal{G}}$ is an equivalence relation.

For $\sigma \in \Sigma_k$, we denote its equivalence class under $\sim_{\mathcal{G}}$ by $[\sigma]_{\mathcal{G}} = \{\hat{\sigma} \in \Sigma_k : \hat{\sigma} \sim_{\mathcal{G}} \sigma\}$ and the quotient space (set of all equivalence classes) by $\Sigma_k / \sim_{\mathcal{G}}$. This set, as shown in Section 4.3, is closely related to the “generically” distinct elements in $\mathcal{P}(A_1, \dots, A_k)$ required by the non-commutativity relations encoded in \mathcal{G} . From the combinatorial point of view, this set is interesting by itself because it is in one-to-one correspondence with the set of acyclic orientations of \mathcal{G} , as the following result shows.

Theorem 4.9. *Let $\mathcal{G} = (V, E)$ be a graph with $V = \{1, 2, \dots, k\}$. Let us define the map*

$$\begin{aligned} \Omega_{\mathcal{G}} : \Sigma_k / \sim_{\mathcal{G}} &\longrightarrow \mathcal{A}(\mathcal{G}) \\ [\sigma]_{\mathcal{G}} &\longmapsto \Omega_{\mathcal{G}}(\sigma), \end{aligned}$$

where, for each $\{i, j\} \in E$, $\Omega_{\mathcal{G}}(\sigma)$ is the orientation given by

$$\Omega_{\mathcal{G}}(\sigma)(\{i, j\}) := (i, j), \quad \text{if } \sigma(i) < \sigma(j). \quad (4.1)$$

Then $\Omega_{\mathcal{G}}$ is well defined (i.e., $\sigma \sim_{\mathcal{G}} \hat{\sigma}$ implies $\Omega_{\mathcal{G}}(\hat{\sigma}) = \Omega_{\mathcal{G}}(\sigma)$), and it is a bijection.

Proof. Let us first show that $\Omega_{\mathcal{G}}$ is well defined. It suffices to show that it is well defined for a single allowed swap $\hat{\sigma} = \tau \circ \sigma$, where $\tau = (i \ i+1)$ and $\{\sigma^{-1}(i), \sigma^{-1}(i+1)\} \notin E$. Since $\tau = (i \ i+1)$, $i_1 < i_2$ implies either $\tau(i_1) < \tau(i_2)$ or $i_1 = i_2 - 1 = i$. Hence, $\sigma(i_1) < \sigma(i_2)$ implies $\tau \circ \sigma(i_1) < \tau \circ \sigma(i_2)$ for all $\{i_1, i_2\} \in E$ since $\{\sigma^{-1}(i), \sigma^{-1}(i+1)\} \notin E$.

Now, let us show that $\Omega_{\mathcal{G}}$ is surjective. Let $\omega \in \mathcal{A}(\mathcal{G})$ be an orientation of \mathcal{G} . According to Proposition 4.6, there is a topological sort that produces a total order \preceq on V , which we write as $i_1 \prec i_2 \prec \dots \prec i_k$. Then we take $\sigma \in \Sigma_k$ defined by $\sigma(j) = i_j$, for $j = 1, \dots, k$, so that $\omega = \Omega_{\mathcal{G}}(\sigma)$.

To prove that $\Omega_{\mathcal{G}}$ is injective, let σ and $\hat{\sigma}$ be two permutations in Σ_k such that $\Omega_{\mathcal{G}}(\hat{\sigma}) = \Omega_{\mathcal{G}}(\sigma)$. Without loss of generality, and relabeling the vertices of V if necessary, we assume that $\sigma = \mathbf{id}$ is the identity permutation. We construct a sequence of permutations, $\sigma_1 = \hat{\sigma}, \sigma_2, \dots, \sigma_k$, by the recurrence

$$\sigma_{i+1} = (i \ i+1 \ \dots \ \sigma_i(i)) \circ \sigma_i, \quad (4.2)$$

for $1 \leq i < k$ (note that this recurrence implicitly requires showing $\sigma_i(i) \geq i$). The proof reduces to the following two claims:

- (i) $\sigma_i^{-1}(j) = j$, for $j < i$ (in particular $\sigma_k = \mathbf{id}$), and
- (ii) $\sigma_{i+1} \sim_{\mathcal{G}} \sigma_i$.

Note that the first claim also implies both $\sigma_i(j) \geq i$ and $\sigma_i^{-1}(j) \geq i$ when $j \geq i$, which justifies the requirement that $\sigma_i(i) \geq i$.

We proceed by induction on i . For $i = 1$, both claims are trivial. We now assume both (i) and (ii) are true up to some $i < k$. If it happens that $\sigma_i(i) = i$, then both claims are trivially satisfied at $i + 1$ with $\sigma_{i+1} = \sigma_i$. Otherwise, for σ_{i+1} as in (4.2), we have, for the first claim:

$$\begin{aligned} \text{For } j < i: \quad \sigma_{i+1}^{-1}(j) &= \sigma_i^{-1} \circ (\sigma_i(i) \ \dots \ i)(j) \\ &= \sigma_i^{-1}(j) = j \\ \text{For } j = i: \quad \sigma_{i+1}^{-1}(i) &= \sigma_i^{-1} \circ (\sigma_i(i) \ \dots \ i)(i) \\ &= \sigma_i^{-1}(\sigma_i(i)) = i. \end{aligned}$$

For the second claim, let $\sigma_{i,j} = (j \ \dots \ \sigma_i(i)) \circ \sigma_i$, for $i \leq j < \sigma_i(i)$. Note that $\sigma_{i,j+1}^{-1}(j) = \sigma_i^{-1}(j)$ and $\sigma_{i,j+1}^{-1}(j+1) = \sigma_i^{-1}(\sigma_i(i)) = i$. Then, $\sigma_{i,j} = (j \ j+1) \circ \sigma_{i,j+1}$ is an allowed swap unless $\{\sigma_i^{-1}(j), i\} \in E$. But we have that $\sigma_i(\sigma_i^{-1}(j)) = j < \sigma_i(i)$ and, on the other hand, by claim (i), we have $\sigma_i^{-1}(j) \geq i$. Hence $\{\sigma_i^{-1}(j), i\} \notin E$, since $\Omega_{\mathcal{G}}(\sigma_i) = \Omega_{\mathcal{G}}(\hat{\sigma}) = \Omega_{\mathcal{G}}(\mathbf{id})$, by the induction and the initial hypotheses. \square

According to Definition 4.1, we introduce the following notion.

Definition 4.10. Given two classes \mathcal{C}_1 and \mathcal{C}_2 in $\Sigma_k / \sim_{\mathcal{G}}$, we say that \mathcal{C}_1 and \mathcal{C}_2 are cyclically related if there are some $\sigma_1 \in \mathcal{C}_1$ and $\sigma_2 \in \mathcal{C}_2$ such that σ_1 and σ_2 are cyclically related.

We note that, unlike the relation for permutations in Σ_k introduced in Definition 4.1, the relation on equivalence classes in Definition 4.10 is not necessarily transitive.

We show in Theorem 4.12 that the cyclic relation of a pair of elements of $\Sigma_k / \sim_{\mathcal{G}}$ corresponds to the following relation between the corresponding elements of $\mathcal{A}(\mathcal{G})$.

Definition 4.11. Let $\mathcal{G} = (V, E)$ be a graph and $\mathcal{A}(\mathcal{G})$ be the set of acyclic orientations of \mathcal{G} . Given $\omega_1, \omega_2 \in \mathcal{A}(\mathcal{G})$, we say that ω_2 is a cut-flip of ω_1 if there is a cut $V = V_1 \cup V_2$ such that, for each edge $\{i, j\} \in E$, we have

- (a) $\omega_1(\{i, j\}) = \omega_2(\{i, j\})$ if $i, j \in V_1$ or if $i, j \in V_2$.
- (b) $\omega_1(\{i, j\}) = (i, j)$ and $\omega_2(\{i, j\}) = (j, i)$ if $i \in V_1$ and $j \in V_2$.

We say that $\omega, \widehat{\omega} \in \mathcal{A}(\mathcal{G})$ are connected by d cut-flips if there exists a sequence $\omega_0 = \omega, \omega_1, \dots, \omega_d = \widehat{\omega} \in \mathcal{A}(\mathcal{G})$, such that ω_{i-1} is a cut-flip of ω_i , for $i = 1, \dots, d$.

We say that $\mathcal{A}(\mathcal{G})$ is connected by d cut-flips if any two orientations are connected by d cut-flips. For $d = 1$ we just say that $\mathcal{A}(\mathcal{G})$ is connected by cut-flips.

We note that, by swapping V_1 and V_2 , the relations in Definition 4.11 are symmetric. In plain language, a cut-flip is a cut where the edges of \mathcal{G} that cross the cut are oriented from V_1 to V_2 in ω_1 and from V_2 to V_1 in ω_2 , while the non-crossing edges of \mathcal{G} have the same orientation in both ω_1 and ω_2 . Theorem 4.12 shows that cut-flips graphically represent the cyclic relations of quotient space $\Sigma_k / \sim_{\mathcal{G}}$.

Theorem 4.12. Let $\mathcal{G} = (V, E)$ be a graph, with $V = \{1, 2, \dots, k\}$, and let $\Omega_{\mathcal{G}}$ be the map defined in Theorem 4.9. Then $[\sigma_1]_{\mathcal{G}}$ and $[\sigma_2]_{\mathcal{G}}$ are cyclically related if and only if $\Omega_{\mathcal{G}}(\sigma_2)$ is a cut-flip of $\Omega_{\mathcal{G}}(\sigma_1)$.

Proof. For brevity, throughout the proof we set $\tau := (1 \ 2 \ \dots \ k)^\ell$. For $\sigma \in \Sigma_k$, we have

$$\tau \circ \sigma(i) = \begin{cases} \sigma(i) + \ell, & \text{if } \sigma(i) \leq k - \ell \\ \sigma(i) + \ell - k, & \text{if } \sigma(i) > k - \ell, \end{cases}$$

and hence

$$\tau \circ \sigma(j) - \tau \circ \sigma(i) = \begin{cases} \sigma(j) - \sigma(i), & \text{if } \sigma(i), \sigma(j) \leq k - \ell \\ < 0, & \text{if } \sigma(i) \leq k - \ell < \sigma(j) \\ > 0, & \text{if } \sigma(j) \leq k - \ell < \sigma(i) \\ \sigma(j) - \sigma(i), & \text{if } \sigma(i), \sigma(j) > k - \ell \end{cases}. \quad (4.3)$$

Now, let us prove the ‘‘only if’’ part of the statement. Let $\sigma \in [\sigma_1]_{\mathcal{G}}$ and $\widehat{\sigma} \in [\sigma_2]_{\mathcal{G}}$ be such that $\widehat{\sigma} = \tau \circ \sigma$, for some ℓ . Recall that, by Theorem 4.9, $\Omega_{\mathcal{G}}(\sigma) =$

$\Omega_{\mathcal{G}}(\sigma_1)$ and $\Omega_{\mathcal{G}}(\widehat{\sigma}) = \Omega_{\mathcal{G}}(\sigma_2)$. Let us consider $V_1 = \{\sigma^{-1}(1), \dots, \sigma^{-1}(k - \ell)\}$ and $V_2 = \{\sigma^{-1}(k - \ell + 1), \dots, \sigma^{-1}(k)\}$.

We now verify that $\Omega_{\mathcal{G}}(\widehat{\sigma})$ is a cut-flip of $\Omega_{\mathcal{G}}(\sigma)$. We analyze all possible situations for an edge $\{i, j\} \in E$:

- $i, j \in V_1$: Then, $\sigma(i), \sigma(j) \leq k - \ell$, by the definition of V_1 . Hence, by (4.3), $\widehat{\sigma}(j) - \widehat{\sigma}(i) = \sigma(j) - \sigma(i)$, so $\Omega_{\mathcal{G}}(\widehat{\sigma})(\{i, j\}) = \Omega_{\mathcal{G}}(\sigma)(\{i, j\})$.
- $i, j \in V_2$: Then, $\sigma(i), \sigma(j) > k - \ell$, by the definition of V_1 . Again, by (4.3), $\widehat{\sigma}(j) - \widehat{\sigma}(i) = \sigma(j) - \sigma(i)$, so $\Omega_{\mathcal{G}}(\widehat{\sigma})(\{i, j\}) = \Omega_{\mathcal{G}}(\sigma)(\{i, j\})$ as well.
- $i \in V_1, j \in V_2$: Then $\sigma(i) \leq k - \ell < \sigma(j)$, by the definition of V_1 and V_2 , so $\sigma(i) < \sigma(j)$. Also, by (4.3), $\widehat{\sigma}(j) - \widehat{\sigma}(i) < 0$, so $\Omega_{\mathcal{G}}(\widehat{\sigma})(\{i, j\}) = (j, i)$, whereas $\Omega_{\mathcal{G}}(\sigma)(\{i, j\}) = (i, j)$.
- $i \in V_2, j \in V_1$: In this case, $\sigma(j) \leq k - \ell < \sigma(i)$, so $\sigma(j) < \sigma(i)$. Again, by (4.3), $\widehat{\sigma}(j) - \widehat{\sigma}(i) > 0$, so $\Omega_{\mathcal{G}}(\widehat{\sigma})(\{i, j\}) = (i, j)$, whereas $\Omega_{\mathcal{G}}(\sigma)(\{i, j\}) = (j, i)$.

Let us now prove the ‘‘if’’ part. Suppose $\Omega_{\mathcal{G}}(\sigma_1)$ is a cut-flip of $\Omega_{\mathcal{G}}(\sigma_2)$ via the cut $V_1 \cup V_2$. Set $\ell := k - |V_1|$. Let \preceq be the total order of V defined by: $i \preceq j$ if and only if $\sigma_1(i) \leq \sigma_1(j)$. Define a permutation $\pi \in \Sigma_k$ by having $\pi^{-1}(1), \dots, \pi^{-1}(k - \ell)$ be the elements of V_1 sorted according to \preceq and $\pi^{-1}(k - \ell + 1), \dots, \pi^{-1}(k)$ be the elements of V_2 also sorted according to \preceq . Note that $\pi(i) \leq k - \ell < \pi(j)$ for all $i \in V_1, j \in V_2$.

From Theorem 4.9, the problem now reduces to showing

$$\Omega_{\mathcal{G}}(\pi) = \Omega_{\mathcal{G}}(\sigma_1) \quad \text{and} \quad \Omega_{\mathcal{G}}(\tau \circ \pi) = \Omega_{\mathcal{G}}(\sigma_2),$$

or equivalently, for all $\{i, j\} \in E$, $\pi(i) < \pi(j)$ if and only if $\sigma_1(i) < \sigma_1(j)$, and $\tau \circ \pi(i) < \tau \circ \pi(j)$ if and only if $\sigma_2(i) < \sigma_2(j)$. As before, we consider the separate cases:

- $i, j \in V_1$: Then $\pi(i), \pi(j) \leq k - \ell$, and, by the definition of π , $\pi(i) < \pi(j)$ if and only if $i < j$, which in turn holds if and only if $\sigma_1(i) < \sigma_1(j)$, by the definition of \preceq . Also, by (4.3), $\tau \circ \pi(j) - \tau \circ \pi(i) = \pi(j) - \pi(i)$. Note that, by the definition of cut-flip, the sign of this difference is in turn equal to $\sigma_2(j) - \sigma_2(i)$.
- $i, j \in V_2$: Similar arguments lead to $\pi(i) < \pi(j)$ if and only if $\sigma_1(i) < \sigma_1(j)$ and $\tau \circ \pi(i) < \tau \circ \pi(j)$ if and only if $\sigma_2(i) < \sigma_2(j)$ also in this case.
- $i \in V_1, j \in V_2$: In this case, we have $\pi(i) < \pi(j)$ by construction and $\sigma_1(i) < \sigma_1(j)$ by hypothesis. Also, $\tau \circ \pi(i) > \tau \circ \pi(j)$ by (4.3) and, by the definition of cut-flip, $\sigma_2(i) > \sigma_2(j)$.
- $i \in V_2, j \in V_1$: With similar arguments, we have $\pi(i) > \pi(j)$, $\sigma_1(i) > \sigma_1(j)$, and $\tau \circ \pi(i) < \tau \circ \pi(j)$, $\sigma_2(i) < \sigma_2(j)$.

□

The main result in this section is Theorem 4.13, which gives a simple characterization of those graphs \mathcal{G} such that $\Sigma_k / \sim_{\mathcal{G}}$ is connected by cut-flips, and, when this is the case, establishes the maximum number of cut-flips needed to connect any two classes.

Theorem 4.13. *Let $\mathcal{G} = (V, E)$ be a graph with k vertices. Then $\mathcal{A}(\mathcal{G})$ is connected by cut-flips if and only if \mathcal{G} is a forest. Furthermore, if \mathcal{G} is a forest, and $d < k$ is the length of the longest path in \mathcal{G} then any two classes in $\mathcal{A}(\mathcal{G})$ can be connected by no more than d cut-flips.*

Proof. Proceeding by contradiction, we assume that $\mathcal{A}(\mathcal{G})$ is connected by cut-flips and that there is a cycle in \mathcal{G} , given by $\{i_0, i_1\}, \{i_1, i_2\}, \dots, \{i_{m-1}, i_0\} \in E$. Let ω_1 and ω_2 be two acyclic orientations of \mathcal{G} related by a cut-flip. Every cut must be crossed by an even number of edges in the cycle, so the number of edges of the cycle on which ω_1 and ω_2 disagree (that is, the number of j for which $\omega_1(\{i_{j-1}, i_j\}) \neq \omega_2(\{i_{j-1}, i_j\})$) must be even. Hence, any two acyclic orientations with an odd number of disagreeing edges on the cycle cannot be connected by a sequence of cut-flips. Since there always exist two such acyclic orientations, we get a contradiction.

To prove the converse implication, let us assume that \mathcal{G} is a forest. We will prove, by induction on d (the length of the longest path in \mathcal{G}), that \mathcal{G} is connected by, at most, d cut-flips.

For $d = 0$, $\mathcal{A}(\mathcal{G})$ has only one element, since no edges means that there is only one (vacuous) orientation. For $d = 1$, given two different acyclic orientations $\omega, \hat{\omega}$ of \mathcal{G} , we take V_1 to be the set of all i where $\{i, j\} \in E$ for some j and $\omega(\{i, j\}) = (i, j) \neq \hat{\omega}(\{i, j\}) = (j, i)$ with $V_2 = V - V_1$. By this construction, all edges in E where ω and $\hat{\omega}$ agree join vertices which are both in V_2 and each edge where they disagree is oriented from V_1 to V_2 by ω and from V_2 to V_1 by $\hat{\omega}$.

We now assume the result for d . Let L be the set of leaf nodes of \mathcal{G} . Consider the graph $\tilde{\mathcal{G}} = (V - L, \tilde{E})$ obtained from \mathcal{G} by removing L from V with a subset of the edges $\tilde{E} = \{\{i, j\} \in E : i, j \in V - L\}$. $\tilde{\mathcal{G}}$ is a forest with longest path length $d - 2$ (since any maximal path in \mathcal{G} must start and end on leaf nodes). Let ω and $\hat{\omega}$ be two different acyclic orientations of \mathcal{G} . Then ω and $\hat{\omega}$, as functions restricted to \tilde{E} , are orientations of $\tilde{\mathcal{G}}$. By induction, $\omega|_{\tilde{E}}$ and $\hat{\omega}|_{\tilde{E}}$ are connected by at most $d - 2$ cut-flips, $\tilde{\omega}_0 = \omega|_{\tilde{E}}, \tilde{\omega}_1, \dots, \tilde{\omega}_q = \hat{\omega}|_{\tilde{E}}$ for $q \leq d - 2$. For each $\tilde{\omega}_p$ in the sequence, we define an orientation ω_p of \mathcal{G} by extending $\tilde{\omega}_i$ to a function on E taking $\omega_p|_{E - \tilde{E}} = \omega|_{E - \tilde{E}}$ (in other words, ω_p agrees with ω on the edges of the leaf nodes). As $\tilde{\omega}_{i-1}, \tilde{\omega}_i$ are related by a cut-flip with cut \tilde{V}_1, \tilde{V}_2 (where $\tilde{V}_1 \cup \tilde{V}_2 = V - L$), ω_{i-1}, ω_i are related by a cut-flip with cut $V_1 = \tilde{V}_1 \cup L_1, V_2 = \tilde{V}_2 \cup L_2$ where L_1 is the set of all leaf nodes whose parents are in V_1 and $L_2 = L - L_1$.

Hence, ω_q is connected to ω by $q \leq d - 2$ cut-flips, and $\omega_q|_{\tilde{E}} = \hat{\omega}|_{\tilde{E}}, \omega_q|_{E - \tilde{E}} = \omega|_{E - \tilde{E}}$. What remains is to connect ω_q to $\hat{\omega}$ with 2 cut-flips. We identify two disjoint subsets of L corresponding to whether ω_q orients the leaf

first or second in its edge:

$$\begin{aligned}\widehat{M}_1 &= \{i \in L : \{i, j\} \in E, \omega_q(\{i, j\}) = (i, j) \neq \widehat{\omega}(\{i, j\}) = (j, i)\} \\ \widehat{M}_2 &= \{i \in L : \{i, j\} \in E, \omega_q(\{i, j\}) = (j, i) \neq \widehat{\omega}(\{i, j\}) = (i, j)\}.\end{aligned}$$

The cut-flip given by $\widehat{M}_1, V - \widehat{M}_1$ followed by the one given by $\widehat{M}_2, V - \widehat{M}_2$, connects ω_q to $\widehat{\omega}$. \square

By Theorem 4.9, via the map $\Omega_{\mathcal{G}}$ we may identify the quotient space $\Sigma_k / \sim_{\mathcal{G}}$ with the set of acyclic orientations of \mathcal{G} . Theorem 4.12 tells us that $\mathcal{A}(\mathcal{G})$ is connected by cut-flips if and only if any two equivalence classes in $\Sigma_k / \sim_{\mathcal{G}}$ are cyclically related. As a consequence, Theorem 4.13 says that any two classes in $\Sigma_k / \sim_{\mathcal{G}}$ are cyclically related if and only if \mathcal{G} is a forest.

4.3. Commutativity conditions and distance of Segré characteristics

As shown in Proposition 4.4, when there is more than one equivalence class in $\mathcal{P}(A_1, \dots, A_k)$ under \sim_C , it is pointless to ask about the change in the JCF of different permuted products, since the spectrum can be completely different. On the other hand, when there is only one equivalence class, all permuted products have the same nonzero eigenvalues with the same Segré characteristic.

Motivated by Fiedler matrices, we impose certain commutativity conditions such that any two elements of $\mathcal{P}(A_1, \dots, A_k)$ are connected by a sequence of Flanders bridges, as we did in Section 3 for three matrices. Under these conditions, we also analyze the change in the Segré characteristic of the eigenvalue zero for different permuted products. We say that two products Π and $\tilde{\Pi}$ in $\mathcal{P}(A_1, \dots, A_k)$ are related by a sequence of Flanders bridges if there are some $\Pi_1, \dots, \Pi_{d+1} \in \mathcal{P}(A_1, \dots, A_k)$ such that $\Pi_1 = \Pi$, $\Pi_{d+1} = \tilde{\Pi}$ and (Π_i, Π_{i+1}) is a Flanders pair, for $i = 1, \dots, d$.

Our commutativity conditions are encapsulated in the associated graph. More precisely, we are interested in the graph comprising the non-commutativity relations.

Definition 4.14. *Given k matrices $A_1, \dots, A_k \in \mathbb{C}^{n \times n}$, the graph of non-commutativity relations of A_1, \dots, A_k is the graph $\mathcal{G} = (V, E)$ with $V = \{1, 2, \dots, k\}$, such that $\{i, j\} \in E$ if and only if $A_i A_j \neq A_j A_i$, for all $1 \leq i, j \leq k$ with $i \neq j$.*

Given k matrices $A_1, \dots, A_k \in \mathbb{C}^{n \times n}$, the set of permuted products $\mathcal{P}(A_1, \dots, A_k)$ can be analyzed in the light of the combinatorial approach of Section 4.2. In particular, if $\mathcal{G} = (V, E)$ is the graph of non-commutativity relations of A_1, \dots, A_k , and $\sigma \in \Sigma_k$, then an allowed swap of σ exchanges two consecutive factors in Π_{σ} , $A_{\sigma^{-1}(i)} A_{\sigma^{-1}(i+1)} \mapsto A_{\sigma^{-1}(i+1)} A_{\sigma^{-1}(i)}$. This is allowed because $\{\sigma^{-1}(i), \sigma^{-1}(i+1)\} \notin E$, which means that the matrices $A_{\sigma^{-1}(i)}$ and $A_{\sigma^{-1}(i+1)}$ commute. Hence, the equivalence classes in the quotient space $\Sigma_k / \sim_{\mathcal{G}}$ correspond to the generically distinct elements of $\mathcal{P}(A_1, \dots, A_k)$ obtained by the commutativity relations of the complementary graph of \mathcal{G} . Here the word “generic” means that, for some particular A_1, \dots, A_k , it may happen that some permuted products

in $\mathcal{P}(A_1, \dots, A_k)$ coincide even if they belong to different equivalence classes. For instance, $ABC = ACB$ is possible even if $BC \neq CB$, though this is not generically the case.

Hence, if we consider the elements of $\mathcal{P}(A_1, \dots, A_k)$ as “formal products”, that is, as words of k letters, A_1, \dots, A_k , then $[\sigma_1]_{\mathcal{G}}$ and $[\sigma_2]_{\mathcal{G}}$ are cyclically related if and only if there is a Flanders bridge between Π_{σ_1} and Π_{σ_2} . Then, using Theorem 4.12, the first part of Theorem 4.13 can be stated as follows.

Theorem 4.15. *Let $A_1, \dots, A_k \in \mathbb{C}^{n \times n}$ and \mathcal{G} be the graph of non-commutativity relations of A_1, \dots, A_k . Let permuted products in $\mathcal{P}(A_1, \dots, A_k)$ be considered as formal products (that is, as words of k letters, A_1, \dots, A_k). Then any two products in $\mathcal{P}(A_1, \dots, A_k)$ are related by a sequence of Flanders bridges if and only if \mathcal{G} is a forest.*

Now, from Theorem 4.15 and part (i) of Corollary 2.3, we conclude that when \mathcal{G} is a forest, all permuted products in $\mathcal{P}(A_1, \dots, A_k)$ have the same nonzero eigenvalues together with their corresponding Segré characteristics. The remaining question is to analyze what happens to the zero eigenvalue in this case. Theorem 4.16, which is the main result in this section, establishes an upper bound for the distance of the Segré characteristic at zero of two permuted products, and shows that the bound is attainable. This bound comes from the number of Flanders bridges in the sequence that relates two arbitrary permuted products Π_{σ_1} and Π_{σ_2} ; that is, the number of cut-flips connecting $\Omega_{\mathcal{G}}(\sigma_1)$ and $\Omega_{\mathcal{G}}(\sigma_2)$.

Theorem 4.16. *Let $A_1, \dots, A_k \in \mathbb{C}^{n \times n}$ and \mathcal{G} be the graph of non-commutativity relations of A_1, \dots, A_k . Assume that \mathcal{G} is a forest and let d be the length of the longest path in \mathcal{G} . Then, given $\Pi_1, \Pi_2 \in \mathcal{P}(A_1, \dots, A_k)$, we have*

$$\|\mathcal{S}_0(\Pi_1) - \mathcal{S}_0(\Pi_2)\|_{\infty} \leq d. \quad (4.4)$$

Moreover, this bound is attainable in the following sense: Let \mathcal{G} be any forest with k vertices, and let $d \leq k$ be the length of the longest path in \mathcal{G} . Then there exist k matrices A_1, \dots, A_k such that \mathcal{G} is the graph of non-commutativity relations of A_1, \dots, A_k , and there are $\Pi_1, \Pi_2 \in \mathcal{P}(A_1, \dots, A_k)$ with

$$\|\mathcal{S}_0(\Pi_1) - \mathcal{S}_0(\Pi_2)\|_{\infty} = d. \quad (4.5)$$

Proof. The first part of the statement is an immediate consequence of Theorem 4.13. More precisely, let \mathcal{G} , Π_1 , and Π_2 be as in the statement. By Theorem 4.13, $\mathcal{A}(\mathcal{G})$ is connected by at most d cut-flips. This implies, using Theorem 4.12, that there are $\sigma_1, \sigma_2, \dots, \sigma_{d+1} \in \Sigma_k$ such that $\Pi_1 = \Pi_{\sigma_1}$, $\Pi_{\sigma_{d+1}} = \Pi_2$, and such that $(\Pi_{\sigma_i}, \Pi_{\sigma_{i+1}})$ are Flanders pairs, for $i = 1, \dots, d$. Now, Theorem 1.1 gives

$$\|\mathcal{S}_0(\Pi_{\sigma_i}) - \mathcal{S}_0(\Pi_{\sigma_{i+1}})\|_{\infty} \leq 1, \quad \text{for } i = 1, \dots, d,$$

so we have

$$\|\mathcal{S}_0(\Pi_1) - \mathcal{S}_0(\Pi_2)\|_{\infty} \leq \sum_{i=1}^d \|\mathcal{S}_0(\Pi_{\sigma_i}) - \mathcal{S}_0(\Pi_{\sigma_{i+1}})\|_{\infty} \leq d.$$

For the second part of the statement (attainability of the bound (4.4)), we first consider the case where \mathcal{G} is a path of length d , and define the $(d+1) \times (d+1)$ matrices

$$\begin{aligned}\tilde{A}_i &= \text{diag}(I_{d-i}, J_2(0), I_{i-1}), \quad \text{for } i = 1, \dots, d, \\ \tilde{A}_{d+1} &= \text{diag}(0, I_d).\end{aligned}\tag{4.6}$$

The graph of non-commutativity relations of $\tilde{A}_1, \dots, \tilde{A}_{d+1}$ is a single path of length d from \tilde{A}_1 to \tilde{A}_{d+1} . Moreover, we have $\Pi_1 = \tilde{A}_1 \tilde{A}_2 \cdots \tilde{A}_{d+1} = J_{d+1}(0)$, and $\Pi_2 = \tilde{A}_{d+1} \cdots \tilde{A}_2 \tilde{A}_1 = 0_{d+1}$, so $\|\mathcal{S}_0(\Pi_1) - \mathcal{S}_0(\Pi_2)\|_\infty = d$.


If $\mathcal{G} = (V, E)$ is a tree with $V = \{1, \dots, k\}$, let us assume, without loss of generality, that $\{1, 2\}, \{2, 3\}, \dots, \{d, d+1\}$ is a path of length d in \mathcal{G} . Now, let $\tilde{A}_1, \dots, \tilde{A}_{d+1}$ be as in (4.6), and $\tilde{A}_{d+2} = \cdots = \tilde{A}_k = I_{d+1}$. Let us number the edges in \mathcal{G} different from $\{1, 2\}, \{2, 3\}, \dots, \{d, d+1\}$, as e_1, \dots, e_g . For each of these edges we build up the following k matrices: for the edge $e_s = \{i, j\}$, with $1 \leq s \leq g$, let $D_1^{(s)}, \dots, D_k^{(s)}$ be k nonsingular matrices of size 2×2 such that $D_i^{(s)} D_j^{(s)} \neq D_j^{(s)} D_i^{(s)}$, and $D_\ell^{(s)} = I_2$ for $\ell \neq i, j$. Now, set $A_i = \text{diag}(\tilde{A}_i, D_i^{(1)}, \dots, D_i^{(g)})$, for $i = 1, \dots, k$. The graph of non-commutativity relations of A_1, \dots, A_k is \mathcal{G} , by construction. Moreover, since $D_i^{(s)}$ is nonsingular, for all $s = 1, \dots, g$ and $i = 1, \dots, k$, we have $\Pi_1 := A_1 A_2 \cdots A_k = \text{diag}(J_{d+1}(0), M_1)$ and $\Pi_2 := A_k \cdots A_2 A_1 = \text{diag}(0_{d+1}, M_2)$, with M_1, M_2 nonsingular, so $\mathcal{S}_0(\Pi_1) = (d+1)$, and $\mathcal{S}_0(\Pi_2) = (1, \dots, 1)$ (containing $d+1$ ones), hence $\|\mathcal{S}_0(\Pi_1) - \mathcal{S}_0(\Pi_2)\|_\infty = d$.

Finally, let \mathcal{G} be a forest with t trees. Let k_1, \dots, k_t be the number of vertices in each tree, with $k_1 + \cdots + k_t = k$, and let d_1, \dots, d_t be the lengths of the longest path in each tree. By hypothesis, we have $\max\{d_j : j = 1, \dots, t\} = d$. For each tree, say the j th one, we define matrices $A_1^{(j)}, \dots, A_{k_j}^{(j)} \in \mathbb{C}^{n_j \times n_j}$ as before, such that the graph of non-commutativity of $A_1^{(j)}, \dots, A_{k_j}^{(j)}$ is precisely this tree, and such that $\|\mathcal{S}_0(A_1^{(j)} A_2^{(j)} \cdots A_{k_j}^{(j)}) - \mathcal{S}_0(A_{k_j}^{(j)} \cdots A_2^{(j)} A_1^{(j)})\|_\infty = d_j$. Now, we set $A_i = \text{diag}(\hat{A}_1^{(i)}, \dots, \hat{A}_t^{(i)})$, for $i = 1, \dots, k$, where $\hat{A}_j^{(i)} = A_h^{(j)}$, if $i = k_1 + \cdots + k_{j-1} + h$, for some $1 \leq h \leq k_j$ (where we set $k_0 := 0$), and $\hat{A}_j^{(i)} = I_{n_j}$ otherwise. For these matrices, we have $\|\mathcal{S}_0(A_1 A_2 \cdots A_k) - \mathcal{S}_0(A_k \cdots A_2 A_1)\|_\infty = \max_{j=1, \dots, t} \|\mathcal{S}_0(A_1^{(j)} A_2^{(j)} \cdots A_{k_j}^{(j)}) - \mathcal{S}_0(A_{k_j}^{(j)} \cdots A_2^{(j)} A_1^{(j)})\|_\infty = d$, $\|\mathcal{S}_0(A_1 A_2 \cdots A_k) - \mathcal{S}_0(A_k \cdots A_2 A_1)\|_\infty = \max_{j=1, \dots, t} d_j = d$, and the graph of non-commutativity relations of A_1, \dots, A_k is \mathcal{G} , by construction. \square

The construction in the proof of Theorem 4.16 does not necessarily give the minimum size of A_1, \dots, A_k that satisfy the second part of the statement. Also note that $d \leq k-1$, and equality holds if and only if \mathcal{G} is a path of length $k-1$.

We mention that, in the case in the Fiedler matrices M_1, \dots, M_n [3] the graph of non-commutativity relations is a forest. Moreover, it is just a path from M_1 to M_n .

Example 4.17. Let $\mathcal{G} = (V, E)$, with $V = \{1, 2, \dots, 9\}$, be the following graph:



The length of the longest path in \mathcal{G} is $d = 4$, which corresponds, for instance, to the path $\{9, 1\}, \{1, 3\}, \{3, 8\}, \{8, 7\}$.

Now, let us follow the construction in the second part of the proof of Theorem 4.16 to get 9 particular matrices A_1, \dots, A_9 such that \mathcal{G} is the graph of non-commutativity relations of A_1, \dots, A_9 (by identifying the j th vertex of \mathcal{G} with the matrix A_j) and so that the products $\Pi_1 = (A_9 A_1 A_3 A_8 A_7) A_6 A_2 A_5 A_4$ and $\Pi_2 = (A_7 A_8 A_3 A_1 A_9) A_6 A_2 A_5 A_4$ satisfy $\|\mathcal{S}_0(\Pi_1) - \mathcal{S}_0(\Pi_2)\|_\infty = 4$.

Set $\tilde{A}_9 = \text{diag}(I_3, J_2(0))$, $\tilde{A}_1 = \text{diag}(I_2, J_2(0), 1)$, $\tilde{A}_3 = \text{diag}(1, J_2(0), I_2)$, $\tilde{A}_8 = \text{diag}(J_2(0), I_3)$, $\tilde{A}_7 = \text{diag}(0, I_4)$, and $\tilde{A}_i = I_5$, for $i \neq 1, 3, 7, 8, 9$. Now, let us number (and label) the edges that are not in the path $\{9, 1\}, \{1, 3\}, \{3, 8\}, \{8, 7\}$ as follows: $e_1 = \{3, 6\}, e_2 = \{2, 3\}, e_3 = \{3, 5\}, e_4 = \{4, 5\}$, and set:

$$\begin{aligned} A_1 &= \text{diag}(\tilde{A}_1, I_8), & A_2 &= \text{diag}(I_7, D_2^{(2)}, I_4), \\ A_3 &= \text{diag}(\tilde{A}_3, D_3^{(1)}, D_3^{(2)}, D_3^{(3)}, I_2), & A_4 &= \text{diag}(I_{11}, D_4^{(4)}), \\ A_5 &= \text{diag}(I_9, D_5^{(3)}, D_5^{(4)}), & A_6 &= \text{diag}(I_5, D_6^{(1)}, I_6), \\ A_7 &= \text{diag}(\tilde{A}_7, D_2^{(2)}, I_4), & A_8 &= \text{diag}(\tilde{A}_8, I_8), & A_9 &= (\tilde{A}_9, I_8), \end{aligned}$$

with $D_j^{(i)}$ being nonsingular 2×2 matrices such that $D_3^{(1)} D_6^{(1)} \neq D_6^{(1)} D_3^{(1)}$, $D_3^{(2)} D_2^{(2)} \neq D_2^{(2)} D_3^{(2)}$, $D_3^{(3)} D_5^{(3)} \neq D_5^{(3)} D_3^{(3)}$, and $D_4^{(4)} D_5^{(4)} \neq D_5^{(4)} D_4^{(4)}$. Under these conditions, the graph of non-commutativity relations of A_1, \dots, A_9 is \mathcal{G} , and we have

$$\Pi_1 = (A_9 A_1 A_3 A_8 A_7) A_6 A_2 A_5 A_4 = \text{diag}(J_5(0), J),$$

and

$$\Pi_2 = (A_7 A_8 A_3 A_1 A_9) A_6 A_2 A_5 A_4 = \text{diag}(0_5, J),$$

with $J = \text{diag}(D_3^{(1)} D_6^{(1)}, D_3^{(2)} D_2^{(2)}, D_3^{(3)} D_5^{(3)}, D_5^{(4)} D_4^{(4)})$. Since the matrices $D_j^{(i)}$ are nonsingular, we have that $\mathcal{S}_0(\Pi_1) = (5)$ and $\mathcal{S}_0(\Pi_2) = (1, 1, 1, 1, 1)$, so $\|\mathcal{S}_0(\Pi_1) - \mathcal{S}_0(\Pi_2)\|_\infty = 4$.

Theorem 4.13 implies that if the graph of non-commutativity relations of A_1, \dots, A_k has no cycles, then all permuted products in $\mathcal{P}(A_1, \dots, A_k)$ have the same eigenvalues with their corresponding Segré characteristics. The reverse implication, however, is not true. For example, take $A_1, \dots, A_k \in \mathbb{C}^{n \times n}$ to be upper triangular and such that no pair commutes and the products of the (i, i) diagonal entries of all matrices, $\pi_i = A_1(i, i) A_2(i, i) \cdots A_k(i, i)$, satisfy $\pi_i \neq \pi_j$ for $i \neq j$. Then, all permuted products have the same eigenvalues, with the same Segré characteristic (they are all simple eigenvalues). However, the graph of non-commutativity relations is the complete graph with k vertices, which is far from a forest.

5. Conclusions and open problems

In this paper we have analyzed the change in the JCF of products of k square matrices under permutations of the factors. As an immediate consequence of a classical result by Flanders, the products are classified into equivalence classes under cyclic permutations of the factors, in such a way that the structure in the JCF for nonzero eigenvalues coincide in any two products belonging to the same class. We have first shown that, if no assumptions are imposed on the factors, then any two products belonging to different classes under cyclic permutations may have different nonzero eigenvalues. Moreover, for three matrices, we have seen that it is always possible to prescribe the nonzero eigenvalues of ABC and ACB , with the condition that the product of all eigenvalues coincide for both products. However, we have seen that this prescription is not always possible for more than three matrices.

We have further shown that, by imposing certain commutativity conditions on the factors, the structure in the JCF corresponding to nonzero eigenvalues coincides for all products. In particular, we have seen that this is always the case if the graph of non-commutativity relations of the factors is a forest. We proved that, when considering formal products, there is only one equivalence class, up to cyclic permutations of the factors, if and only if the graph of non-commutativity relations of the factors is a forest. When this graph is a forest, we obtained an upper bound on the difference between the structure (sizes of Jordan blocks) of the JCF associated with the eigenvalue zero in any pair of products, and we saw that this bound is attainable. Moreover, in the case of three matrices, we proved that it is always possible to prescribe the sizes of the blocks associated with zero in the JCF of ABC and ACB as long as the difference between the corresponding sizes is at most 2.

We conclude with some open problems that arise as a natural continuation of the problems addressed in this paper.

- **Open problem 1:** Is it always possible to prescribe the n eigenvalues of the $(k-1)!$ classes under cyclic permutations, provided that the product of all eigenvalues is the same, for k, n satisfying $(k-1)!(n-1) + 1 \leq kn^2$ and $k \geq 4$?
- **Open problem 2:** Given $d \geq 4$ and two nondecreasing sequences $\boldsymbol{\mu}, \boldsymbol{\mu}'$ of nonnegative integers such that $\|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_\infty \leq d-1$, is it always possible to find d matrices, A_1, \dots, A_d , such that their graph of non-commutativity relations is a path, and such that $\mathcal{S}_0(A_1 \cdots A_d) = \boldsymbol{\mu}$ and $\mathcal{S}_0(A_d \cdots A_1) = \boldsymbol{\mu}'$? (The extension of Theorem 4.15 to $d \geq 4$).
- **Open problem 3:** If $M, Q \in \mathbb{C}^{n \times n}$ are such that $\mathcal{S}_\lambda(M) = \mathcal{S}_\lambda(Q)$, for all $\lambda \neq 0$, and $\|\mathcal{S}_0(M) - \mathcal{S}_0(Q)\|_\infty \leq 2$, are there three matrices $A, B, C \in \mathbb{C}^{n \times n}$ with $AC = CA$, such that $M = ABC$ and $Q = CBA$?
- **Open problem 4:** Obtain necessary and sufficient conditions for all products of a given set of k matrices to have the same nonzero eigenval-

ues and corresponding Segré characteristics (in the notation of the paper: $\mathcal{S}_\lambda(\Pi_1) = \mathcal{S}_\lambda(\Pi_2)$, for all $\lambda \neq 0$, and for all $\Pi_1, \Pi_2 \in \mathcal{P}(A_1, \dots, A_k)$).

Acknowledgments: We thank Enrique Arrondo for pointing out reference [15, Th. 7, Ch. I §6], Richard A. Brualdi for recommending the LaTeX packages TikZ and PGF, and the handling Editor of the manuscript, R. A. Horn, for his constructive suggestions and warm support. We also acknowledge useful comments on the exposition by James Hook and a fruitful discussion with Gil Strang.

- [1] S. Bernau and A. Abian. Jordan canonical forms of matrices AB and BA . *Rend. Inst. Mat. Univ. Trieste*, 20(1):101–108, 1988.
- [2] B. Bollobás. *Modern Graph Theory*. Springer, New York, 1998.
- [3] M. Fiedler. A note on companion matrices. *Linear Algebra Appl.*, 372:325–332, 2003.
- [4] H. Flanders. Elementary divisors of AB and BA . *Proc. Amer. Math. Soc.*, 2(6):871–874, 1951.
- [5] S. Furtado and C. R. Johnson. Order invariant spectral properties for several matrices. *Linear Algebra Appl.*, 432:1950–1960, 2010.
- [6] J. Gelonch and C. R. Johnson. Generalization of Flanders’ theorem to matrix triples. *Linear Algebra Appl.*, 380:151–171, 2004.
- [7] J. Gelonch, C. R. Johnson, and P. Rubi3. An extension of Flanders theorem to several matrices. *Linear Multilinear Algebra*, 43:181–200, 1997.
- [8] R. A. Horn, C. R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, 1985.
- [9] R. A. Horn, C. R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, New York, 1991.
- [10] R. A. Horn and D. I. Merino. Contragredient Equivalence: A canonical form and some applications. *Linear Algebra Appl.*, 214:43–92, 1995.
- [11] C. R. Johnson and E. S. Schreiner. The relationship between AB and BA . *Amer. Math. Monthly*, 103(7):578–581, 1996.
- [12] R. A. Lippert and G. Strang. The Jordan forms of AB and BA . *Electron. J. Linear Algebra*, 18:281–288, 2009.
- [13] J. A. McHugh. *Algorithmic Graph Theory*. Prentice-Hall, New Jersey, 1990.
- [14] W. V. Parker and B. E. Mitchell. Elementary divisors of certain matrices. *Duke Mathematical Journal*, 19(3):483–485, 1952.
- [15] I. R. Safarevich. *Basic Algebraic Geometry*. Springer, Berlin, 1974.
- [16] R. C. Thompson. On the matrices AB and BA . *Linear Algebra Appl.*, 1(1):43 – 58, 1968.