

Decomposing Sets of Inversions

Lukas Katthän*

Fachbereich Mathematik und Informatik
Philipps-Universität Marburg
Germany

katthaen@mathematik.uni-marburg.de

Submitted: Aug 8, 2012; Accepted: Feb 14, 2013; Published: Feb 25, 2013

Mathematics Subject Classifications: 05A05, 52B12, 05E40

Abstract

In this paper we consider the question how the set of inversions of a permutation $\pi \in S_n$ can be partitioned into two subsets, which are themselves inversion sets of permutations in S_n . Our method is to study the modular decomposition of the inversion graph of π . A correspondence to the substitution decomposition of π is also given. Moreover, we consider the special case of multiplicative decompositions.

Keywords: Inversion sets; Permutation graphs; Simple Permutations; Linear ordering polytope

1 Introduction

For a permutation $\pi \in S_n$ denote its *inversion set* by

$$T(\pi) := \{ \{i, j\} \in \mathbb{N}^2 \mid 1 \leq i < j \leq n, \pi(i) > \pi(j) \} .$$

In this paper, we address the following problem:

Problem 1.1. For a given permutation $\pi \in S_n$, give a description of all $\tau_1, \tau_2 \in S_n$ such that

$$\begin{aligned} T(\tau_1) \cup T(\tau_2) &= T(\pi) \\ T(\tau_1) \cap T(\tau_2) &= \emptyset \\ \tau_1, \tau_2 &\neq id_n . \end{aligned} \tag{1}$$

*This work was partially supported by the DAAD

In other words, we want to find all ways to distribute the inversions of π into two disjoint sets, such that each is itself the inversion set of a permutation. The motivation behind this problem is as follows. In [12], toric statistical ranking models are considered. One of these models is the *inversion model*, which is also known as *Babington-Smith Model* in the statistics literature, see [9]. The toric ideal I_{BS} associated to this model is the kernel of the map

$$k[X_\pi \mid \pi \in S_n] \rightarrow k[X_{ij} \mid 1 \leq i < j \leq n]$$

$$X_\pi \mapsto \prod_{\{i,j\} \in T(\pi)} X_{ij}$$

It follows from general theory that I_{BS} is generated by differences of monomials [11, Lemma 4.1]. By construction, a difference of monomials $\prod_i X_{\pi_i} - \prod_i X_{\tau_i}$ is contained in I_{BS} if and only if $\bigcup_i T(\pi_i) = \bigcup_i T(\tau_i)$ as multisets. Thus the generators of the ideal I_{BS} encode the relations among the inversion sets of permutations. Therefore, a set of generators for this ideal not only provides algebraic information for the Babington-Smith Model but also encodes fundamental information about the combinatorics of permutations. However, I_{BS} turns out to be a rather large and complex object, for example the authors of [12] found computationally that for $n = 6$ there are as many as 130377 quadratic generators and there are also generators of higher degree. Therefore, as a first step in understanding this object, we study its quadratic generators for all n . The ideal I_{BS} is invariant under the right action of the S_n , so if $m := X_{\pi_1} X_{\pi_2} - X_{\tau_1} X_{\tau_2} \in I_{BS}$, then also $m\pi_1^{-1} = X_{id_n} X_{\pi_2 \pi_1^{-1}} - X_{\tau_1 \pi_1^{-1}} X_{\tau_2 \pi_1^{-1}} \in I_{BS}$. Therefore we can restrict our attention to binomials of the form

$$X_{id_n} X_\pi - X_{\tau_1} X_{\tau_2}.$$

From our discussion, the following observation is immediate:

Proposition 1.2. *A binomial $X_{id_n} X_\pi - X_{\tau_1} X_{\tau_2}$ lies in I_{BS} if and only if π, τ_1 and τ_2 satisfy (1).*

Thus Problem 1.1 is equivalent to the problem of describing the quadratic generators of I_{BS} . In the recent preprint [6], the following closely related question is considered: Let $\omega_{0,n} \in S_n$ denote the permutation of maximal length (i.e. the one mapping $i \mapsto n+1-i$).

Problem 1.3. Give a description of all sets $\{\tau_1, \dots, \tau_l\} \subset S_n$ such that $T(\omega_{0,n}) = \bigcup_i T(\tau_i)$ and $T(\tau_i) \cap T(\tau_j) = \emptyset$ for $i \neq j$.

The motivation and the methods employed by the authors of [6] are different from ours, but some intermediate results of this paper were also found independently there. In particular, Proposition 4.1 and part of Theorem 3.9 resemble Proposition 2.2 and Proposition 3.14 in [6].

Another perspective on a toric model is via its *model polytope*. The model polytope associated to the inversion model is the *linear ordering polytope* [12], which is a well-studied object in combinatorial optimization, see [10, Chapter 6]. In [14] the following question is addressed:

Problem 1.4. Which permutations $\pi \in S_n$ are neighbours of the identity permutation in the graph of the linear ordering polytope?¹

In [14], a characterization of these permutations is obtained, but as we show after Theorem 3.9 there is a gap in the proof. Nevertheless, the result from [14] is correct and we extend the result and provide a proof in Theorem 3.9. It turns out that a permutation has a decomposition as in (1) if and only if it is not a neighbour of the identity permutation in the graph of the linear ordering polytope. However, in the present paper we are interested in a description of all possible decompositions of type (1).

This paper is divided into four sections and an appendix. In Section 2 we review the concept of modular decomposition for graphs, the characterisation of inversion sets of permutations and we discuss blocks of permutations. In Section 3, we prove our main results. In Theorem 3.3, we give an answer to Problem 1.1 in terms of the modular decomposition of the inversion graph of π . Moreover, we consider a modification of (1), where we impose the further restriction that $\pi = \tau_1\tau_2$. We show in Theorem 3.6 that if π admits a solution of (1), then it also admits a solution satisfying $\pi = \tau_1\tau_2$. Since Problem 1.1 is formulated without referring to graphs, in Theorem 3.13 we give a reformulation of Theorem 3.3 which avoids notions from graph theory. In Section 4, we show that the problem of decomposing an inversion set into three or more inversion sets can be reduced to (1). Moreover, we show that permutations of sufficiently high length always admit a solution of (1). In the appendix we prove a result connecting the blocks of a permutation with the modules of its inversion graph. The result from the appendix seems rather natural to us, but since we were not able to find it in the literature, we include a proof.

2 Preliminaries

2.1 Notation

Let us first fix some notation. We denote a graph G on a vertex set V with edge set $E \subset V \times V$ by $G = (V, E)$. All our graphs are undirected and simple. For two vertices v, w , let vw denote the (undirected) edge between v and w . We say v and w are connected in G if $vw \in E$ and we write $vw \in G$ by abuse of notation.

For a natural number $n \in \mathbb{N}$, we write $[n]$ for the set $\{1, \dots, n\}$. For a finite set S , we write $\binom{S}{2}$ for the set of subsets of S containing exactly 2 elements. For two sets A and B , we denote by $A \dot{\cup} B$ the disjoint union of A and B . For $\pi \in S_n$ we denote by $T(\pi)$ the *inversion set*

$$\left\{ \{i, j\} \in \binom{[n]}{2} \mid i < j, \pi(i) > \pi(j) \right\}.$$

This set can be considered as the edge set of an undirected graph $G(\pi) = ([n], T(\pi))$, the *inversion graph* of π . We consider this graph without the natural order on its vertices, therefore in general $G(\pi)$ does not uniquely determine π . The graphs arising this way

¹The linear ordering polytope is called the “permutation polytope” in [14].

are called *permutation graphs*, see [4]. By another abuse of notation, we write $ij \in T(\pi)$ (resp. $ij \in G(\pi)$) if $\{i, j\}$ is an inversion of π . For two subsets $A, B \subset [n]$, we write $A < B$ if $a < b$ for every $a \in A, b \in B$.

2.2 Modular decomposition of graphs

In this subsection we review the modular composition for graphs, see [4, Chapter 1.5] for a reference. Let $G = (V, E)$ be a graph.

Definition 2.1 ([4]). 1. A set $M \subset V$ is called a *module* of G if for $m_1, m_2 \in M$ and $v \in V \setminus M$ it holds that $vm_1 \in E$ if and only if $vm_2 \in E$.

2. A module M is called *strong* if for every other module N either $M \cap N = \emptyset$, $M \subset N$ or $N \subset M$ holds.

In [4, p. 14] it is shown that for every module there is a unique minimal strong module containing it. A graph is called *prime* if V and its vertices are its only modules. We denote by $\overline{G}$ the complementary graph $\overline{G} = (V, \binom{V}{2} \setminus E)$ of G . For a subset $U \subset V$, we denote by G_U the induced subgraph of G on U .

Theorem 2.2 (Theorem 1.5.1, [4]). *Let $G = (V, E)$ be a graph with at least two vertices. Then the maximal strong submodules (m.s.s.) of G form a partition of V and exactly one of the following conditions hold:*

Parallel case G is not connected. Then its m.s.s. are its connected components.

Serial case $\overline{G}$ is not connected. Then the m.s.s. of G are the connected components of $\overline{G}$.

Prime case Both G and $\overline{G}$ are connected. Then there is a subset $U \subset V$ such that

1. $\#U > 3$,
2. G_U is a maximal prime subgraph of G ,
3. and every m.s.s. M of G has $\#M \cap U = 1$.

We call a module M of G *parallel*, *serial* or *prime* corresponding to which condition of the above theorem is satisfied by G_M . As a convention, we consider single vertices as parallel modules. By the following lemma, we do not need to distinguish between modules of G contained in a module M and modules of G_M .

Lemma 2.3. *Let G be a graph, M a module of G and $U \subset M$ a subset. Then U is a module of G if and only if it is a module of G_M . Moreover, U is a strong module of G if and only if it is strong as a module of G_M .*

Proof. The first statement is immediate from the definitions. For the second statement, first assume that U is not strong as a module in G_M . We say that a module N overlaps U if $N \cap U \neq \emptyset$, $N \not\subseteq U$ and $U \not\subseteq N$ holds. So by our assumption, there is a module $N \subset M$ of G_M overlapping U . But N is also a module of G , hence U is not strong as a module of G . On the other hand, if U is not strong as a module of G , then there is a module N of G overlapping U . Now, $M \setminus N$ is a module of G ([4, Prop 1.5.1 (ii)]), and thus a module of G_M . But $M \setminus N$ overlaps U , so U is not strong as a module of G_M . $\square$

If M and N are two disjoint modules of G , then one of the following holds:

1. Either every vertex of M is connected to every vertex of N . Then we call M and N connected in G and we write MN for the set of edges between vertices of M and N .
2. Otherwise no vertex of M is connected to any vertex of N .

The edges connecting the m.s.s. of a module M are called *external edges* of M . So M is parallel if and only if it has no external edges. Note that every edge of G is an external edge for exactly one strong module. We close this section by giving a description of the non-strong modules of G :

Lemma 2.4. *Let G be a graph and let M be a module which is not strong. Then M is the union of some m.s.s. of a parallel or serial strong module. On the other hand, any union of m.s.s. of a parallel or serial strong module is a module.*

Proof. Let N be the smallest strong module containing M . The m.s.s. of N partition it, so M is a union of some of them. If N is prime, then consider the set U in Theorem 2.2. Since M is not strong, it is a union of at least two but not of all m.s.s. of N . So $M \cap U$ is a nontrivial submodule of G_U , contradicting Theorem 2.2. Hence N is either serial or parallel.

For the converse, let M be a union of m.s.s. of a serial or parallel strong module N . By Lemma 2.3, it suffices to prove that M is a module of G_N . Let $x, y \in M$ and $m \in N \setminus M$. The edges xm, ym are both external in N . But if N is serial, it has all possible external edges and if it is parallel, it has none at all. In both cases, the claim is immediate. $\square$

2.3 Inversion sets and blocks

We recall the characterization of those sets that can arise as inversion sets of a permutation.

Proposition 2.5 (Proposition 2.2 in [13], see also [3]). *Let $T \subset \binom{[n]}{2}$ be a subset. The following conditions are equivalent:*

1. *There exists a permutation $\pi \in S_n$ with $T = T(\pi)$.*
2. *For every $1 \leq i < j < k \leq n$ it holds that:*

- If $ij, jk \in T$, then $ik \in T$.
- If $ik \in T$, then at least one of ij and jk lies in T .

If a subset $T \subset \binom{[n]}{2}$ satisfies the conditions of above proposition, say $T = T(\pi)$, then so does its complement by $\binom{[n]}{2} \setminus T = T(\omega_{0,n}\pi)$. We now take a closer look at the modules of the inversion graph of a permutation $\pi \in S_n$. Let us call a set $I \subset [n]$ of consecutive integers an *interval*.

Definition 2.6 ([5]). 1. A π -block is an interval $I \subset [n]$ such that its image $\pi(I)$ is again an interval.

2. A π -block is called *strong* if for every other π -block J either $I \cap J = \emptyset$, $I \subset J$ or $J \subset I$ holds.

The importance of π -blocks for our purpose stems from the following theorem:

Theorem 2.7. Let $I \subset [n]$ and $\pi \in S_n$. The following implications hold:

1. I is a π -block $\implies I$ is a module of $G(\pi)$
2. I is a strong π -block $\iff I$ is a strong module of $G(\pi)$

In particular, every strong module of $G(\pi)$ is an interval.

The first part of this theorem is relatively easy to prove and is mentioned in [14]. Its converse fails for trivial reasons: By Lemma 2.4, the non-strong modules of $G(\pi)$ are exactly the unions of m.s.s. of parallel or serial strong modules of $G(\pi)$. But such a union is not necessarily an interval. A complete proof of Theorem 2.7 is included in the appendix. We call a π -block parallel, serial or prime if it is a module of this type.

3 Main results

In this section we prove our main results. Fix a permutation $\pi \in S_n$. For $\tau_1, \tau_2 \in S_n$, we will write $\pi = \tau_1 \sqcup \tau_2$ to indicate that the three permutations satisfy (1). We call $\tau_1 \sqcup \tau_2$ an *inv-decomposition* of π . If an inv-decomposition of π exists, we call π *inv-decomposable*.

3.1 Inversion decomposition

In this subsection, we describe all possible inv-decompositions of π . We start with an elementary observation:

Lemma 3.1. Let $i, j, k \in [n]$ such that $ij, ik \in G(\pi)$ and $jk \notin G(\pi)$. Assume that $\pi = \tau_1 \sqcup \tau_2$ for $\tau_1, \tau_2 \in S_n$. Then ij, ik are both either in $G(\tau_1)$ or in $G(\tau_2)$.

Proof. We consider the different relative orders of i, j and k separately, but we may assume $j < k$.

$i < j < k$ The edge ik is contained either in $T(\tau_1)$ or in $T(\tau_2)$, say in $T(\tau_1)$. By assumption $jk \notin T(\tau_1)$, therefore by Proposition 2.5 we have $ij \in T(\tau_1)$.

$j < i < k$ This case is excluded by Proposition 2.5.

$j < k < i$ Analogous to the first case.

□

Note that there is no assumption on the relative order of i, j and k , so this is really a statement about the inversion graph of π . Lemma 3.1 gives rise to a partition of the edges of $G(\pi)$: Two edges $ij, ik \in G(\pi)$ with a common endpoint are in the same *edge class* if $jk \notin G(\pi)$, and our partition is the transitive closure of this relation. Thus by Lemma 3.1 two edges in the same class always stay together when we distribute the inversions of π on τ_1 and τ_2 . In [7] edge classes are considered for a different motivation. In that paper the following description is given².

Proposition 3.2 ([7]). *Let $G = (V, E)$ be a graph with at least two vertices. Then there are two kinds of edge classes:*

1. *For two m.s.s. $M_1, M_2 \subset M$ of a serial module M , the set M_1M_2 is an edge class.*
2. *The set of external edges of a prime module forms an edge class.*

Every edge class is of one of the above types.

Edge classes are also considered in [8, Chapter 5] under the name “colour classes” and in [14] as the connected components of a certain graph Γ_π . Theorem 1 in the latter reference gives a different characterization of edge classes. Now we can state our main result. We give a description of all ways of partitioning $T(\pi)$ into two sets satisfying (1).

Theorem 3.3. *Consider a partition $T(\pi) = T_1 \dot{\cup} T_2$ of the inversion set of π into non-empty subsets $T_1, T_2 \subset T(\pi)$. For such a partition, the following conditions are equivalent:*

1. *There exist permutations $\tau_1, \tau_2 \in S_n$ such that $T_i = T(\tau_i)$ for $i = 1, 2$. In particular, $\pi = \tau_1 \sqcup \tau_2$.*
2. *For every strong prime module of $G(\pi)$, all its external edges are either in T_1 or in T_2 . For every strong serial module of $G(\pi)$ with p maximal strong submodules $M_1 < \dots < M_p$ there exists a permutation $\sigma \in S_p$, such that for each pair $1 \leq i < j \leq p$ it holds that $M_iM_j \subset T_1$ if and only if $ij \in T(\sigma)$.*

Proof. (1) $\Rightarrow$ (2): Every edge of $G(\pi)$ is an external edge of a module M that is either prime or serial. If M is a prime module, then its external edges form an edge class, hence they all are in T_1 or T_2 . If M is a serial module with m.s.s. $M_1, \dots, M_p$, then the sets M_iM_j are edge classes. For every M_i , choose a representative $a_i \in M_i$. We construct a

²Note that what we call module is called “geschlossene Menge” (closed set) in [7].

permutation $\sigma \in S_p$ as follows: Order the images $\tau_1(a_i), i = 1, \dots, p$ in the natural order. Then $\sigma(i)$ is the position of $\tau_1(a_i)$ in this order. Thus, for $i < j$ we have

$$\begin{aligned} M_i M_j \subset T(\tau_1) &\iff \tau_1(a_i) > \tau_1(a_j) \\ &\iff \sigma(i) > \sigma(j) \\ &\iff ij \in T(\sigma) \end{aligned}$$

(2) $\Rightarrow$ (1): By symmetry, we only need to show the existence of τ_1 . For this, we verify conditions of Proposition 2.5. This is a condition for every three numbers $1 \leq i < j < k \leq n$, so let us fix them. Note that our hypothesis on T_1 and T_2 implies that every edge class of $T(\pi)$ is contained either in T_1 or T_2 .

Let M be the smallest strong module containing these three numbers. It holds that i and k are in different m.s.s. of M , because every strong module containing both would also contain j , since it is an interval by Theorem 2.7. Now we distinguish two cases: Either, i and j are in the same m.s.s. of M , or all three numbers are in different m.s.s..

In the first case, let $i, j \in M_a$ and $k \in M_b$. Then $ik, jk \in M_a M_b$ belong to the same edge class, so either both or neither of them are in T_1 . This is sufficient to prove that the criterion is satisfied.

In the second case, the edges ij, jk, ik are all external to M . Hence, if M is prime, either none of them is in T_1 or all that are also in $T(\pi)$. Since $T(\pi)$ is the inversion set of a permutation, the criterion of Proposition 2.5 is clearly satisfied in this case. If M is serial, then the edges correspond to inversions of σ : Let $i \in M_a, j \in M_b, k \in M_c$, then $M_a M_b \subset T_1$ if and only if $ab \in T(\sigma)$ and similarly for the other edges. Since σ is a permutation, the criterion is again satisfied. $\square$

As a corollary, we can count the number of inv-decompositions of π :

Corollary 3.4. *Let m be the number of strong prime modules and let k_i be the number of strong serial modules with i maximal strong submodules, $2 \leq i \leq n$. The number of inv-decompositions of π is*

$$\frac{1}{2} 2^m \prod_{i=2}^n (i!)^{k_i} - 1$$

In particular, the number of inv-decompositions depends only on the inversion graph $G(\pi)$.

We exclude the trivial inv-decomposition $\pi = \pi \sqcup id_n$, therefore the “ -1 ” in above formula. The factor $\frac{1}{2}$ is there because we identify $\tau_1 \sqcup \tau_2 = \tau_2 \sqcup \tau_1$.

3.2 Multiplicative decompositions

A notable special case of an inv-decomposition is the following:

Definition 3.5. We call an inv-decomposition $\pi = \tau_1 \sqcup \tau_2$ *multiplicative* if $\pi = \tau_1 \tau_2$ or $\pi = \tau_2 \tau_1$ (multiplication as permutations).

This kind of inv-decomposition is surprisingly common. In this subsection, we prove the following

Theorem 3.6. *Every inv-decomposable permutation has a multiplicative inv-decomposition. Moreover, if a permutation π has a non-multiplicative inv-decomposition and $G(\pi)$ is connected, then π has a decreasing subsequence of size 4.*

The assumption that $G(\pi)$ is connected is needed to avoid a rather trivial case. $G(\pi)$ is disconnected if and only if π maps a lower interval $[k] \subsetneq [n]$ to itself. So in this case, π is the product of a permutation π_1 on $[k]$ and a permutation π_2 on $\{k+1, \dots, n\}$. If we have multiplicative inv-decompositions $\pi_1 = \tau_{11}\tau_{12}$ and $\pi_2 = \tau_{21}\tau_{22}$, then $\pi = \tau_{11}\tau_{22} \sqcup \tau_{12}\tau_{21}$ is in general not multiplicative. Before we prove Theorem 3.6, we prepare two lemmata.

Lemma 3.7. *If $C \subset [n]$ is the set of vertices of a connected component of $G(\pi)$, then $\pi(C) = C$.*

Proof. Consider $i \in [n] \setminus C$ and $c \in C$. If $i < c$, then $\pi(i) < \pi(c)$ and the same is true for “ $>$ ”, thus the claim follows from bijectivity. $\square$

Lemma 3.8. *Assume $\pi = \tau_1 \sqcup \tau_2$. If every connected component of $G(\tau_2)$ is an induced subgraph of $G(\pi)$, then $\pi = \tau_1\tau_2$.*

Proof. We will prove that $T(\tau_1\tau_2) = T(\pi)$. Let $M_1, \dots, M_s$ be the vertex sets of the connected components of $G(\tau_2)$. Note that $[n] = \bigcup M_k$. By [2, Ex 1.12] it holds that $T(\tau_1\tau_2)$ is the symmetric difference of $T(\tau_2)$ and $\tau_2^{-1}T(\tau_1)\tau_2$. However, we observe that in our situation the two sets are disjoint, thus the symmetric difference is actually a disjoint union. To see this, note that every edge of $G(\tau_2)$ has both endpoints in the same M_k for a $1 \leq k \leq s$, and every edge of $G(\tau_1)$ has its endpoints in different sets. Since by Lemma 3.7 it holds that $\tau_2(M_k) = M_k$ for every k , this property is preserved under the conjugation with τ_2 . Hence, the sets are disjoint.

Next, we prove that every M_k is a $G(\pi)$ -module. So fix a k . If M_k has only one element, then it is trivially a $G(\pi)$ -module, so assume that M_k has more than one element. Let M' be the smallest strong module of $G(\pi)$ containing M_k and let G_k be the subgraph of $G(\pi)$ induced by M_k . Because $\pi = \tau_1 \sqcup \tau_2$ is a valid decomposition, G_k is a union of edge classes. Thus, if M' is prime, we conclude that $M_k = M'$ and we are done. If M' is parallel, then G_k cannot be connected, thus we only need to consider the case that M' is serial. But in this case, M_k is a union of m.s.s. of M' because of the form of the edge classes, given by Proposition 3.2. By Lemma 2.4, we conclude that M_k is indeed a module of $G(\pi)$. Moreover, it follows that M_k is also a module of $G(\tau_1)$, because $G(\pi)$ and $G(\tau_1)$ differ only inside the M_k . Finally, consider the set

$$\tau_2^{-1}T(\tau_1)\tau_2 = \bigcup_{k,l} \{ \{ \tau_2(i), \tau_2(j) \} \mid \{ i, j \} \in T(\tau_1), i \in M_k, j \in M_l \} .$$

Because $\tau_2(M_k) = M_k$ and M_k is a module of $G(\tau_1)$ for all k , it holds that

$$\{ \{ \tau_2(i), \tau_2(j) \} \mid \{ i, j \} \in T(\tau_1), i \in M_k, j \in M_l \} = \{ \{ i, j \} \in T(\tau_1) \mid i \in M_k, j \in M_l \}$$

for all k and l . Hence $\tau_2^{-1}T(\tau_1)\tau_2 = T(\tau_1)$ and the claim follows. $\square$

Proof of Theorem 3.6. For the first statement, assume that π is inv-decomposable. Then by Corollary 3.4 there are either at least two non-parallel strong π -blocks $I_1, I_2 \subset [n]$, or at least one serial strong π -block I_3 with at least three m.s.s..

In the first case, we may assume $I_1 \not\subseteq I_2$. We set T_2 to be the set of edges in the induced subgraph of $G(\pi)$ on I_2 . In the second case, we set T_2 to be the set of edges in the induced subgraph of G on the union of the two first m.s.s. of I . In both cases, we set $T_1 = T(\pi) \setminus T_2$. By Theorem 3.3, this is a valid inv-decomposition, and by Lemma 3.8 it is multiplicative.

For the second statement, we will prove that $G(\pi)$ contains a complete subgraph on 4 vertices. Let $\pi = \tau_1 \sqcup \tau_2$ be a non-multiplicative inv-decomposition. Consider a minimal path from 1 to n in $G(\pi)$. If i and j are two vertices in this path that are not adjacent in this path, then they are not adjacent in $G(\pi)$, because otherwise we had a shortcut. Thus by Lemma 3.1 we conclude that every edge in this path lies in the same edge class. Hence either $G(\tau_1)$ or $G(\tau_2)$ contains a path connecting 1 with n , say $G(\tau_1)$. By Lemma 3.1 this implies that $G(\tau_1)$ has no isolated vertices.

By our hypothesis and by Lemma 3.8, there exists a connected component of $G(\tau_2)$ that is not an induced subgraph of $G(\pi)$. Then there exist $1 \leq i, j \leq n$ such that $ij \in G(\tau_1)$ and there is a minimal path $i, i', \dots, j$ connecting i and j in $G(\tau_2)$. By Lemma 3.1 we have $i'j \in G(\pi)$. We also want to make sure that $i'j \in G(\tau_2)$. If this is not the case, then replace i by i' . Then the corresponding statements still hold, but the minimal path is shorter. Thus, by induction we may assume $i'j \in G(\tau_2)$. Since $G(\tau_1)$ has no isolated vertices, there is a vertex k such that $i'k \in G(\tau_1)$. Again by Lemma 3.1 we conclude that $ik, jk \in G(\pi)$. Thus $G(\pi)$ contains the complete subgraph on i, i', j and k . $\square$

3.3 Characterization of inv-decomposability

We use the results we have proven so far to derive a characterization of inv-decomposability. Let us recall the definition of the Linear Ordering Polytope. To every permutation π we associate a vector $v_\pi \in \mathbb{R}^{n^2}$ by setting

$$(v_\pi)_{ij} = \begin{cases} 1 & \text{if } \pi(i) < \pi(j), \\ 0 & \text{otherwise.} \end{cases}$$

The Linear Ordering Polytope is defined to be the convex hull of these vectors. The inv-decomposability of a permutation π can now be characterized as follows.

Theorem 3.9. *For $\pi \in S_n$ the following statements are equivalent:*

1. *There exist $\tau_1, \tau_2 \in S_n \setminus \{id_n\}$ such that $T(\pi) = T(\tau_1) \dot{\cup} T(\tau_2)$ and $\pi = \tau_1 \tau_2$, i.e. π has a multiplicative inv-decomposition.*
2. *There exist $\tau_1, \tau_2 \in S_n \setminus \{id_n\}$ such that $T(\pi) = T(\tau_1) \dot{\cup} T(\tau_2)$, i.e. π is inv-decomposable.*
3. *v_π is not a neighbour of the identity in the graph of the linear ordering polytope.*

4. There are at least two edge classes of $G(\pi)$.
5. There are at least two (not necessarily strong) non-trivial non-parallel π -blocks. (By a non-trivial π -block, we mean a π -block that is neither a singleton nor $[n]$)

In [14], the implications $(2) \Rightarrow (3) \Rightarrow (4) \Rightarrow (5) \Rightarrow (2)$ are proven, although the condition (2) is not explicitly mentioned. As indicated in Section 1, there is a gap in the proof. Indeed on page 4 of [14], in the proof of the implication $(3) \Rightarrow (4)$ the following argument is used. If v_π is not a neighbour of v_{id_n} , then there is a point on the line between the points that can be written as a convex combination of other vertices, e.g. $\lambda v_{id_n} + (1-\lambda)v_\pi = \sum \lambda_i v_{\tau_i}$ for $\lambda, \lambda_i \in [0, 1]$ and the λ_i sum up to 1. Considering the support set of the vectors on the left and right-hand side of this equation we obtain an expression $T(\pi) = \bigcup T(\tau_i)$. Note, that in general this union is not disjoint. In [14], the existence of this expression, together with the assumption that $G(\pi)$ has only one edge class leads to a contradiction, proving $(3) \Rightarrow (4)$. But $T(2413) = \{13, 23, 24\} = T(2314) \cup T(1423)$ and $G(2413)$ has only one edge class, providing a counterexample to above argument. Since the notation of [14] is different from ours, we provide a full proof of the implications for the convenience of the reader.

Proof. $1 \Leftrightarrow 2$: Theorem 3.6.

$2 \Rightarrow 3$: If $T(\pi) = T(\tau_1) \dot{\cup} T(\tau_2)$, then the midpoint of the line connecting v_{id_n} and v_π is also the midpoint of the line connecting v_{τ_1} and v_{τ_2} , thus it cannot be an edge.

$3 \Rightarrow 4$: If v_π is not a neighbour of v_{id_n} , then we can write $\lambda v_{id_n} + (1-\lambda)v_\pi = \sum \lambda_i v_{\tau_i}$ for $\lambda, \lambda_i \in [0, 1]$ and $\tau_i \neq \pi$ for every i . We clear denominators to make the coefficients integral. The important observation is that every non-zero component of the right-hand side has the same value.

Consider $a, b, c \in [n]$ such that $ab, bc \in T(\pi)$ and $ac \notin T(\pi)$. Then b cannot lie between a and c because of Proposition 2.5. There remain four possible relative orders of a, b and c . We assume $b < a < c$, the other cases follow analogously. Every τ_i with $bc \in T(\tau_i)$ has also $ba \in T(\tau_i)$, again by Proposition 2.5. But the number of τ_i having the inversion bc equals the number of those having ba . Hence, every τ_i has either both or none of the inversions. It follows that if $T(\tau_i)$ contains an inversion, then it already contains the whole edge class of it. Thus if $G(\pi)$ has only one edge class, then for every i either $\tau_i = \pi$ or $\tau_i = id_n$, which is absurd.

$4 \Rightarrow 5$: This follows from the description of the edge classes, Proposition 3.2.

$5 \Rightarrow 2$: Under our hypothesis, the formula in Corollary 3.4 cannot evaluate to zero.

□

3.4 Substitution decomposition

We give a reformulation of Theorem 3.3 avoiding notions from graph theory. For this, we employ the concept of *substitution decomposition*, which was introduced in [1], see [5] for a survey. We start by giving an explicit description of the three types of π -blocks.

Proposition 3.10. *Let $I \subset [n]$ be a π -block with at least two elements and let $I_1 < \dots < I_l$ be its maximal strong submodules.*

1. *I is parallel if and only if $\pi(I_1) < \pi(I_2) < \dots < \pi(I_l)$.*
2. *I is serial if and only if $\pi(I_1) > \pi(I_2) > \dots > \pi(I_l)$.*
3. *Otherwise I is prime.*

Proof. This is a consequence of Theorem 2.7. I is parallel if and only if it has no external edges. This translates to the statement that the relative order of the I_i is preserved. Similarly, I is serial if and only if it has all possible external edges. Again, this translates to the statement that the relative order of the I_i is reversed. $\square$

In the remainder of this section, we consider permutations as words $\pi = \pi_1\pi_2\cdots\pi_n$. The *size* of a permutation is the number of letters in its word³. The special word $id_n := 12\cdots(n-1)n$ is called an *identity*. If $\pi = \pi_1\pi_2\cdots\pi_n$ is a permutation, we call $\tilde{\pi} := \pi_n\pi_{n-1}\cdots\pi_1$ the *reversal* of π . The word $\omega_{0,n} := id_n = n(n-1)\cdots 21$ is called *reverse identity*. Two finite sequences $a_1, \dots, a_q$ and $b_1, \dots, b_q$ of natural numbers are called *order isomorphic* whenever $a_i < a_j$ if and only if $b_i < b_j$. Given a permutation $\pi \in S_m$ and m further permutations $\sigma_1, \dots, \sigma_m$ of not necessarily the same size, we define the *inflation* $\pi[\sigma_1, \dots, \sigma_m]$ by replacing the value $\pi(i)$ by an interval order isomorphic to σ_i . For a more detailed treatment of the inflation operation see [6]. A permutation π is called *simple* if there are no other π -blocks than $[n]$ and the singletons. Note that by Theorem 2.7 a permutation π is simple if and only if its inversion graph $G(\pi)$ is prime.

Proposition 3.11. *Every permutation π can be uniquely expressed as an iterated inflation, such that every permutation appearing in this expression is either an identity, a reverse identity or a simple permutation, and no identity or reverse identity is inflated by a permutation of the same kind.*

We call this the substitution decomposition of π . It is slightly different from the decomposition in [5]. The existence of our decomposition follows from the existence of the decomposition given in that paper, but we consider it to be instructive for our discussion to give a proof nevertheless.

Proof. Let $I_1 < I_2 < \dots < I_l$ be the maximal strong π -subblocks of $[n]$. Define a permutation $\alpha \in S_l$ by requiring $\alpha(i) < \alpha(j) \Leftrightarrow \pi(I_i) < \pi(I_j)$ for $1 \leq i, j \leq l$. Moreover, let σ_i be the permutation order isomorphic to $\pi(I_i)$ for $1 \leq i \leq l$. Then $\pi = \alpha[\sigma_1, \dots, \sigma_l]$. By

³This is called “length” in [5] but we reserve that notion for the number of inversions.

Theorem 2.2 the π -block $[n]$ is either parallel, serial or prime. Hence by Proposition 3.10 we conclude that α is either an identity, a reverse identity or simple. By applying this procedure recursively to the σ_i , we get the claimed decomposition.

The last claim follows also from Theorem 2.2, because it implies that no serial module has a maximal strong submodule which is again serial, and the same for parallel modules. This is just the statement that connected components of a graph are connected. $\square$

The proof gives a correspondence between the strong π -blocks and the permutations appearing in the substitution decomposition. The strong parallel, serial and prime π -blocks correspond to the identities, reverse identities and simple permutations, respectively. Now we can reformulate Theorem 3.3 in terms of inflations:

Construction 3.12. *Let π be a permutation. Define two new permutations τ_1, τ_2 in the following way: Write down two copies of the substitution decomposition of π . For every simple permutation α in it, replace α in one of the copies by an identity. For every reverse identity, replace it in one copy by an arbitrary permutation σ of the same size and in the other by the reverse $\check{\sigma}$. Then let τ_1 and τ_2 be the permutations defined by these iterated inflations.*

Theorem 3.13. *Let π, τ_1, τ_2 be permutations as above and assume that $\tau_1, \tau_2 \neq id_n$. Then $\pi = \tau_1 \sqcup \tau_2$ and every pair (τ_1, τ_2) satisfying this condition can be found this way.*

Proof. This is immediate from Theorem 3.3 using the correspondence described above. $\square$

4 Further results

In this section, we give some further results. First, we consider the generalisation of (1) to more than two components. It turns out that this case can easily be reduced to the case of two components, as the next proposition shows.

Proposition 4.1. *Let $\pi, \tau_1, \dots, \tau_l \in S_n$ be permutations such that $T(\pi) = \bigcup T(\tau_i)$ and $T(\tau_i) \cap T(\tau_j) = \emptyset$ for $i \neq j$. Then for every $1 \leq i, j \leq l$ there exists a $\tau_{ij} \in S_n$ such that $T(\tau_{ij}) = T(\tau_i) \dot{\cup} T(\tau_j)$.*

Proof. We show that $T := T(\tau_i) \cup T(\tau_j)$ satisfies the condition of Proposition 2.5. Fix $1 \leq a_1 < a_2 < a_3 \leq n$. Note that $\binom{[n]}{2} \setminus T = \binom{[n]}{2} \setminus T(\tau_i) \cap \binom{[n]}{2} \setminus T(\tau_j)$, so if $a_1 a_2 \notin T$ and $a_2 a_3 \notin T$, then $a_1 a_3 \notin T$. On the other hand, if $a_1 a_2, a_2 a_3 \in T$, then $a_1 a_3 \in T(\pi)$ and thus $a_1 a_3 \in T(\tau_k)$ for some k . But then $T(\tau_k)$ contains also $a_1 a_2$ or $a_2 a_3$, therefore k equals i or j . It follows that $a_1 a_3 \in T$. $\square$

From Theorem 3.9 we can derive a simple sufficient (but by no means necessary) condition for a permutation to be inv-decomposable.

Proposition 4.2. *Every permutation $\pi \in S_n$ with at least $\binom{n}{2} - n + 2$ inversions is inv-decomposable*

Proof. Let $\tau := \omega_{0,n}\pi$. Then $T(\tau)$ is the complement of $T(\pi)$ and the τ has at most $n - 2$ inversions. So the graph $G(\tau)$ is disconnected, because it has n vertices, but only $n - 2$ edges. This means that $[n]$ is a serial module of $G(\pi)$. If π is not inv-decomposable, then $[n]$ can have only two maximal strong submodules, both parallel. But then $G(\tau)$ would be the disjoint union of two complete graphs. This is not possible with the restriction on the number of edges, as a direct calculation shows. $\square$

5 Appendix: Blocks and modules

In this appendix, we prove the following theorem:

Theorem 5.1. *Let $I \subset [n]$ and $\pi \in S_n$. The following implications hold:*

1. I is a π -block $\implies I$ is a module of $G(\pi)$
2. I is a strong π -block $\iff I$ is a strong module of $G(\pi)$

For the rest of this section, let $\pi \in S_n$ denote a fixed permutation. For brevity, we write block for π -block and modules are to be understood as modules of $G(\pi)$. Recall that a block is an interval whose image under π is again an interval. The first statement of Theorem 5.1 is a direct consequence of the following lemma.

Proposition 5.2. *Let $I \subset [n]$ be an interval. Then I is a module if and only if it is a block.*

Proof.

$$\begin{aligned}
 I \text{ module} &\iff \forall i \in [n] \setminus I : [\exists j \in I : ij \in G(\pi) \Rightarrow \forall j \in I : ij \in G(\pi)] \\
 &\iff \forall i \in [n] \setminus I : [\exists j \in I : \pi(i) < \pi(j) \Rightarrow \forall j \in I : \pi(i) < \pi(j)] \\
 &\iff \forall i \in [n] \setminus I : \pi(i) < \pi(I) \text{ or } \pi(i) > \pi(I) \\
 &\iff I \text{ block}
 \end{aligned}$$

$\square$

We split the proof of the second part of Theorem 5.1 into three lemmata. For a set $S \subset [n]$ we define $S_{<} := \{x \in [n] \mid x < S\}$ and similarly $S_{>}$. We also define $S_{><} := [n] \setminus (S_{<} \cup S \cup S_{>}) = \{x \in [n] \mid \exists a, b \in S : a < x < b, x \notin S\}$.

Lemma 5.3. *Let M be a module. Then $\pi(M_{<} \cup M_{>}) = \pi(M)_{<} \cup \pi(M)_{>}$ and $\pi(M_{><}) = \pi(M)_{><}$.*

Proof. Let i be in $M_{<}$. If $ij \in G(\pi)$ for all $j \in M$, then $\pi(i) \in \pi(M)_{>}$. Otherwise $ij \notin G(\pi)$ for all $j \in M$ and $\pi(i) \in \pi(M)_{<}$. A similar argument for $i \in M_{>}$ proves that $\pi(M_{<} \cup M_{>}) \subset \pi(M)_{<} \cup \pi(M)_{>}$. For $i \in M_{><}$ there exist $j, k \in M$ with $j < i < k$. If $ij \in G(\pi)$, then also $ik \in G(\pi)$ and therefore $\pi(j) > \pi(i) > \pi(k)$. Otherwise $\pi(j) < \pi(i) < \pi(k)$. Hence $\pi(M_{><}) \subset \pi(M)_{><}$. Equality follows for both inclusions because π is bijective. $\square$

Lemma 5.4. *Every strong module is a strong block.*

Proof. Let M be a strong module but not an interval. We write $M \cup M_{><} = M_1 \cup M_2 \cup \dots \cup M_l$ where the M_i are the interval components of M and $M_{><}$ and $M_1 < M_2 < \dots < M_l$. We proceed by proving the following list of claims:

1. $M \cup M_{><}$ is a module.
2. $M_{><}$ is a module.
3. Either $\pi(M_1) < \pi(M_2) < \dots < \pi(M_l)$ or $\pi(M_1) > \pi(M_2) > \dots > \pi(M_l)$.
4. $M_1 \cup M_2$ is a module.

The last claim is a contradiction to the assumption that M is strong, because $M_1 \subset M$ and $M_2 \cap M = \emptyset$. Hence M must be an interval. By Proposition 5.2 we conclude that it is a block. Every other block is also a module, hence the strongness as a block follows from the strongness as a module. We prove the claims one after the other:

1. From Lemma 5.3 we know $\pi(M_{><}) = \pi(M)_{><}$ and hence $\pi(M \cup M_{><}) = \pi(M) \cup \pi(M)_{><}$. Thus this set is a block and the claim follows from Proposition 5.2.
2. Because $M \cup M_{><}$ is a module, by Lemma 2.3 it suffices to prove that $M_{><}$ is a module of $M \cup M_{><}$. Let $i, j \in M_{><}, k \in M$ and $ik \in G(\pi)$. We need to prove $jk \in G(\pi)$. Choose $k_1, k_2 \in M$ such that $k_1 < i, j < k_2$. Because $ik \in G(\pi)$ and M is a module we know that $k_1i, ik_2 \in G(\pi)$. Now we use Proposition 2.5 to conclude:

$$\begin{aligned} k_1i, ik_2 \in G(\pi) &\Rightarrow k_1k_2 \in G(\pi) \\ &\Rightarrow k_1j, jk_2 \in G(\pi) \\ &\Rightarrow jk \in G(\pi) \end{aligned}$$

3. It suffices to prove for every $1 < i < l$: Either $\pi(M_{i-1}) < \pi(M_i) < \pi(M_{i+1})$ holds or the corresponding statement with “ $>$ ” holds. If this is wrong, there are $x_k \in M_k, k \in \{i-1, i, i+1\}$ with $\pi(x_{i-1}) > \pi(x_i) < \pi(x_{i+1})$ or $\pi(x_{i-1}) < \pi(x_i) > \pi(x_{i+1})$, say, the first. But then $x_{i-1}x_i \in G(\pi)$ and $x_ix_{i+1} \notin G(\pi)$. But both edges are in $MM_{><}$, so this is a contradiction to the previous claim.
4. Since $M_1 \cup M_2$ is an interval, by Proposition 5.2 it suffices to prove that $\pi(M_1 \cup M_2)$ is also an interval. For $x \in [n] \setminus (M_1 \cup M_2)$, it holds that either $x \in M_{<} \cup M_{>}$ or $x \in M_3 \cup \dots \cup M_l$. In the first case we know by Lemma 5.3 that $\pi(x) \in \pi(M)_{<} \cup \pi(M)_{>} \subset \pi(M_1 \cup M_2)_{<} \cup \pi(M_1 \cup M_2)_{>}$. For $x \in M_3 \cup \dots \cup M_l$ it follows from the previous claim that $\pi(x) \in \pi(M_1 \cup M_2)_{<} \cup \pi(M_1 \cup M_2)_{>}$. Therefore, $\pi([n] \setminus (M_1 \cup M_2)) \subset \pi(M_1 \cup M_2)_{<} \cup \pi(M_1 \cup M_2)_{>}$. Because $M_1 \cup M_2$ is an interval we can conclude from this that $\pi(M_1 \cup M_2)_{><} = \emptyset$, thus the claim follows.

□

Lemma 5.5. *Every strong block is a strong module.*

Proof. Suppose $I \subset [n]$ is a strong block. By Proposition 5.2 I is a module. Thus it remains to prove that it is strong, so assume the contrary. By Lemma 2.4 it is the union of m.s.s. of a strong module M' . Write $M' = M_1 \cup \dots \cup M_l$, where the M_i are the m.s.s. We have already proven in Lemma 5.4 that they are intervals. Choose two consecutive ones M_i, M_{i+1} such that $M_i \subset I$ and $M_{i+1} \cap I = \emptyset$. Then $M_i \cup M_{i+1}$ is an interval by construction and a module by Lemma 2.4. Therefore it is a block by Proposition 5.2. But this is a contradiction to the hypothesis that I is strong. $\square$

References

- [1] M.H. Albert and M.D. Atkinson, *Simple permutations and pattern restricted permutations*, Discrete Mathematics **300** (2005), no. 13, 1 – 15, [doi:10.1016/j.disc.2005.06.016](https://doi.org/10.1016/j.disc.2005.06.016).
- [2] A. Björner and F. Brenti, *Combinatorics of Coxeter groups*, Springer, 2005.
- [3] A. Björner and M. L. Wachs, *Permutation statistics and linear extensions of posets*, Journal of Combinatorial Theory, Series A **58** (1991), no. 1, 85 – 114.
- [4] A. Brandstädt, J. P. Spinrad, and V. B. Le, *Graph classes: a survey*, SIAM monographs on discrete mathematics and applications, SIAM, 1999.
- [5] R. Brignall, *A survey of simple permutations*, In Permutation Patterns, London Mathematical Society Lecture Notes Series (376), Cambridge University Press, 2010.
- [6] R. Dewji, I. Dimitrov, A. McCabe, M. Roth, D. Wehlau, and J. Wilson, *Decomposing Inversion Sets of Permutations and Applications to Faces of the Littlewood-Richardson Cone*, (2011), [arXiv:1110.5880](https://arxiv.org/abs/1110.5880).
- [7] T. Gallai, *Transitiv orientierbare Graphen*, Acta Mathematica Hungarica **18** (1967).
- [8] M. C. Golumbic, *Algorithmic graph theory and perfect graphs*, Acad. Press, 1980.
- [9] J. I. Marden, *Analyzing and modeling rank data*, Chapman & Hall, 1995.
- [10] R. Martí and G. Reinelt, *The linear ordering problem: Exact and heuristic methods in combinatorial optimization*, Springer, 2011.
- [11] B. Sturmfels, *Gröbner bases and convex polytopes*, Amer. Math. Soc, 1996.
- [12] B. Sturmfels and V. Welker, *Commutative algebra of statistical ranking*, Journal of Algebra **361** (2012), 264 – 286, [doi:10.1016/j.jalgebra.2012.03.028](https://doi.org/10.1016/j.jalgebra.2012.03.028).
- [13] T. Yanagimoto and M. Okamoto, *Partial orderings of permutations and monotonicity of a rank correlation statistic*, Annals of the Institute of Statistical Mathematics **21** (1969).
- [14] H. P. Young, *On permutations and permutation polytopes*, Polyhedral Combinatorics, Mathematical Programming Studies, vol. 8, North-Holland Publishing Company, 1978.