Available online at www.sciencedirect.com
SCIENCE DIRECT®
European Journal of Combinatorics 25 (2004) 657–673
European Journal of Combinatorics
www.elsevier.com/locate/ejc

Stable sets of maximal size in Kneser-type graphs

Benoit Larose[a], Claudia Malvenuto[b]

[a] *Department of Mathematics and Statistics, Concordia University, 1455 De Maisonneuve West, Montreal, Qc, Canada, H3G 1M8*
[b] *Dipartimento di Informatica, Università di Roma "La Sapienza", Via Salaria, 113, 00198 Roma, Italy*

Received 29 September 2003; accepted 23 October 2003

Abstract

We introduce a family of vertex-transitive graphs with specified subgroups of automorphisms which generalise Kneser graphs, powers of complete graphs and Cayley graphs of permutations. We compute the stability ratio for a wide class of these. Under certain conditions we characterise their stable sets of maximal size.

MSC (2000): 05D05; 05C99

Keywords: Stable sets; Kneser graphs; Permutation graphs

1. Introduction

Consider the graph K_n^m whose vertices are all tuples $(x_1, \ldots, x_m)$ with entries in $\{1, 2, \ldots, n\}$ where two tuples are adjacent if they have no entry in common. In [9] Greenwell and Lovász characterise the stable sets of maximal size in this graph: a set of vertices is a maximal stable set if and only if it consists of all those tuples whose ith entry is some fixed value $1 \leq j \leq n$.

Let n, r be positive integers with $n \geq 2r$. The Kneser graph $K(r, n)$ is the graph whose vertices are the r-element subsets of $\{1, 2, \ldots, n\}$, two of them being adjacent if they are disjoint. The Erdős–Ko–Rado Theorem [7] states that the stability ratio of $K(r, n)$ is r/n; furthermore, it follows from the Hilton–Milner inequality [10] (see also [3] and [8] for

E-mail addresses: larose@mathstat.concordia.ca (B. Larose), claudia@dsi.uniromal.it (C. Malvenuto).
URL: http://cicma.mathstat.concordia.ca/faculty/larose/.

0195-6698/$ - see front matter © 2003 Elsevier Ltd. All rights reserved.
doi:10.1016/j.ejc.2003.10.006

simple proofs of this fact) that the stable sets of maximal size in $K(r, n)$ are precisely those families of sets that contain some fixed element.

One may restate this last result in a manner that underlines the resemblance with the first example: indeed, one may view r-subsets of $\{1, 2, \ldots, n\}$ (henceforth denoted by $[n]$) as n-tuples $(x_1, \ldots, x_n)$ with exactly r entries equal to 1 and $n - r$ entries equal to 0, two tuples being adjacent if they do not have a common entry equal to 1. Then a stable set of maximal size consists of all tuples that have the ith coordinate equal to 1, for a fixed i.

We introduce a class of graphs which generalise these two situations. The vertices are tuples with a fixed distribution of occurrences of the symbols, and the adjacency between two strings will indicate whether a symbol from a fixed subset of symbols appears at different positions in the two strings. Let us rephrase this in a more precise way.

Let $b \geq 1$ be an integer (the number of symbols for the strings) and let $d_1 \geq d_2 \geq \cdots \geq d_b \geq 0$ be integers (the distribution of the symbols), with $n = \sum_i d_i \geq 1$ (the length of the strings). Let m denote the largest index i such that $d_i > 0$. Let P be any subgroup of the symmetric group S_b, and let C be a non-empty subset of $[b] = \{1, \ldots, b\}$. We construct a graph $G(P; C; d_1, \ldots, d_b)$ as follows: its vertices are the n-tuples $(a_1, \ldots, a_n) \in [b]^n$ such that there exists a permutation $\sigma \in P$ and a permutation $\tau \in S_n$ for which $(a_1, \ldots, a_n) = (x_{\tau(1)}, \ldots, x_{\tau(n)})$ where $(x_1, \ldots, x_n) = (\sigma(1), \ldots, \sigma(1), \sigma(2), \ldots, \sigma(2), \ldots, \sigma(m), \ldots, \sigma(m))$ and $\sigma(i)$ appears exactly d_i times. Two such tuples $(x_1, \ldots, x_n)$ and $(y_1, \ldots, y_n)$ are *not adjacent* if and only if there exists a coordinate i such that $x_i = y_i \in C$ (the two tuples coincide for a symbol belonging to C).

We shall keep the above notation and terminology throughout this paper. We also define another useful parameter: given the graph $G(P; C; d_1, \ldots, d_b)$ let

$$d = \max\{d_i : i \in C\}.$$

The group S_n acts naturally on the graph $G(P; C; d_1, \ldots, d_b)$ by permuting the entries of the tuples. Furthermore it is easy to see that the correspondence

$$(x_1, \ldots, x_n) \mapsto (\sigma(x_1), \ldots, \sigma(x_n))$$

for each $\sigma \in P$ defines an action of P on the graph as a group of automorphisms, provided the permutations in P preserve the set C, i.e. $\sigma(i) \in C$ for all $i \in C$ and all $\sigma \in P$. *For the remainder of the paper we shall assume that this condition always holds.* Assuming this, one verifies at once that $G(P; C; d_1, \ldots, d_b)$ is vertex-transitive, under the combined actions of S_n and P.

Example 1. Let us choose the sequence of d_i to be 3, 2, 2, 1, 0, 0 so $n = 8$, $b = 6$ and $m = 4$. Let P consist of the following group of permutations:

$$P = \{id, (2\,1), (3\,4\,5), (3\,5\,4), (2\,1)(3\,4\,5), (2\,1)(3\,5\,4)\}.$$

Then the vertices of the graph are all possible permutations of the coordinates of the following tuples

$$(1, 1, 1, 2, 2, 3, 3, 4)$$
$$(2, 2, 2, 1, 1, 3, 3, 4)$$
$$(1, 1, 1, 2, 2, 4, 4, 5)$$
$$(1, 1, 1, 2, 2, 5, 5, 3)$$
$$(2, 2, 2, 1, 1, 4, 4, 5)$$
$$(2, 2, 2, 1, 1, 5, 5, 3).$$

Example 2. Let $1 \le r \le n$ be integers such that $2r \le n$. Then $K(r, n)$, the Kneser graph of r-subsets of $[n]$, is isomorphic to $G(P; C; d_1, d_2)$ where $b = 2$, P is the trivial group, $C = \{2\}$, $d_1 = n - r$ and $d_2 = r$. The isomorphism is given by the correspondence which sends every r-subset X of $[n]$ to the n-tuple $(x_1, \dots, x_n)$ where $x_i = 2$ if $i \in X$ and $x_i = 1$ otherwise.

Example 3. Let $d_i = 1$ for every i, hence $n = b$. Let P be the full symmetric group and $C = [b]$. With this choice of parameters the vertices of $G(S_n; [n]; 1, 1, \dots, 1)$ are the permutations of $[n]$, two permutations σ and τ being adjacent if there is no coordinate i such that $\sigma(i) = \tau(i)$, i.e. $\sigma\tau^{-1}$ is a derangement. We will denote this graph simply by $G(S_n)$. Note that it is the Cayley graph of the group S_n with generating set the derangements.

Example 4. Let $G(S_b)^{(n)}$ denote the graph $G(S_b; [b]; 1, 1, \dots, 1, 0, \dots, 0)$ where 1 appears n times. The vertices of this graph are the injections of $[n]$ into $[b]$; two of these injections α and β are *not* adjacent if $\alpha(i) = \beta(i)$ for some $1 \le i \le n$.

Let $\mathcal{O}$ denote the orbit under P of any element $q \in C$ such that $d_q = d$. We shall prove the following in Section 3.

Theorem. *The independence ratio of the graph $G = G(P; C; d_1, \dots, d_b)$ satisfies*

$$\frac{\alpha(G)}{|G|} \ge \frac{1}{|\mathcal{O}|} \sum_{i \in \mathcal{O}} \frac{d_i}{n}.$$

We shall prove that in various cases this bound is tight (Theorems 3.2 and 3.3). For example if P is trivial, as in the case of Kneser graphs, this ratio is equal to d/n; if $P = S_b$ the ratio is equal to $1/b$, a result already obtained by Deza and Frankl [5] (see also [6]) in the special case of $G(S_n)$.

We shall then be interested in the actual nature of the stable sets of maximal size of the graph $G(P; C; d_1, \dots, d_b)$. Let $1 \le p \le n$ and $1 \le q \le b$. Let I_p^q denote the induced subgraph of $G(P; C; d_1, \dots, d_b)$ that consists of all tuples $(x_1, \dots, x_n)$ such that $x_p = q$. If $q \in C$ this is an independent set. In fact (see Section 3) it is of maximal size precisely if the above bound is tight.

Problem. For which parameters $P, C, d_1, \dots, d_b$ are all the stable sets of maximal size in the graph $G(P; C; d_1, \dots, d_b)$ of the form I_p^q?

We now outline the contents of the paper. In the next section we introduce the terminology and basic results that we shall require in the sequel. In Section 3 we prove the above-mentioned theorem. In Section 4 we prove that the sets I_p^q are indeed the only stable sets of maximal size in the graph $G(S_n)$ (Theorem 4.1).[1] In Section 5 we extend this to the graphs $G(S_b)^{(n)}$ (Theorem 5.1) and then use this to prove the following more general result:

Theorem. *Let $b \geq 5$. If $n/2 \geq d_1$ and there are at least 3 non-zero d_i's, then the sets I_p^q are the only stable sets of maximal size in the graph $G(S_b; C; d_1, \ldots, d_b)$.*

Finally we conclude with some remarks on open questions and further results (Section 6).

2. Preliminaries and terminology

In this paper all graphs are finite, undirected and without loops. We denote the vertex set and the edge set of a graph G by $V(G)$ and $E(G)$ respectively. Let G be a graph. Recall that a set I of vertices of G is called *stable* (or *independent*) if no two vertices in I are adjacent. The *stability (independence) number* of G, denoted by $\alpha(G)$, is the maximum cardinality of a stable set in G. The *stability (independence) ratio* of G is the ratio of the stability number to the number of vertices. If G and H are graphs, a *homomorphism* from G to H is an edge-preserving map from $V(G)$ to $V(H)$, i.e. a function f such that $\{f(g), f(g')\} \in E(H)$ whenever $\{g, g'\} \in E(G)$. A graph G is *vertex-transitive* if the automorphism group of G acts transitively on $V(G)$, i.e. for every x and y in $V(G)$ we can find an automorphism f of G such that $f(x) = y$.

For $n \geq 1$ we shall denote the complete graph on n vertices by K_n.

One of our main tools is the following result, often referred to as the 'no-homomorphism lemma':

Lemma 2.1 ([2]).[2] *Let G and H be graphs such that H is vertex-transitive and there exists a homomorphism $\phi : G \to H$. Then*

$$\frac{\alpha(G)}{|V(G)|} \geq \frac{\alpha(H)}{|V(H)|}. \tag{1}$$

Furthermore, if equality holds in (1), then for any stable set I of cardinality $\alpha(H)$ in H, $\phi^{-1}(I)$ is a stable set of cardinality $\alpha(G)$ in G.

For $0 \leq i, j \leq n-1$, define the circular distance from i to j, denoted by $\partial_n(i, j)$, as the distance from the vertex i to the vertex j in the cycle of length n, that is the minimum between the representative of $(i-j)$ mod n and the representative of $(j-1)$ mod n in $\{0, 1, \ldots, n-1\}$.

Let r, s be positive integers such that $r < s/2$. The *circular graph* $\mathrm{Circ}(r, s)$ is defined as follows: its set of vertices is $\mathbb{Z}_s = \{0, 1, \ldots, s-1\}$, and two vertices u and v are adjacent if $\partial_s(u, v) \geq r$. The neighborhood of a vertex u of $\mathrm{Circ}(r, s)$ is $\{u+r, \ldots, u+s-r\}$.

[1] At the moment of submitting this paper we learned that this was also obtained recently by Cameron and Ku [4].

[2] The second part of the lemma is not explicitly formulated in the paper by Albertson and Collins, but it is implicit in the proof.

The following result is from [13]:

Lemma 2.2. *The stability ratio of* $\mathrm{Circ}(r, s)$ *is* r/s*, and the only stable sets of maximal size of* $\mathrm{Circ}(r, s)$ *are the arcs*

$$\{k, k+1, \ldots, k+r-1\},$$

$k \in V(\mathrm{Circ}(r, s))$.

3. The independence ratio of $G(P; C; d_1, \ldots, d_b)$

Let $q \in C$ be such that $d_q \geq d_i$ for all $i \in C$ (i.e. $d_q = d$) and let $\mathcal{O}$ denote the orbit of q under P. We derive a bound for the stability ratio of G:

Theorem 3.1. *The independence ratio of the graph* $G = G(P; C; d_1, \ldots, d_b)$ *satisfies*

$$\frac{\alpha(G)}{|G|} \geq \frac{1}{|\mathcal{O}|} \sum_{i \in \mathcal{O}} \frac{d_i}{n}.$$

Proof. Let $I = I_1^q$, the set of all tuples in G whose first coordinate is equal to q. Obviously it is a stable set in G: we compute its cardinality. Let H denote the subgroup of P that consists of all permutations σ that satisfy the condition $d_i = d_{\sigma(i)}$ for all $i \in [b]$. For every $\sigma \in P$ let G_σ denote the induced subgraph of G whose vertices are the n-tuples $(x_1, \ldots, x_n)$ for which there exists $\tau \in S_n$ such that

$$(x_{\tau(1)}, \ldots, x_{\tau(n)}) = (\sigma(1), \ldots, \sigma(1), \sigma(2), \ldots, \sigma(2), \ldots, \sigma(m), \ldots, \sigma(m))$$

where $\sigma(i)$ appears exactly d_i times, $1 \leq i \leq m$. The reader may easily verify the following: the graphs G_σ and G_ρ are either equal or disjoint, and $G_\sigma = G_\rho$ if and only if $\rho^{-1}\sigma \in H$. In particular, a counting argument yields that for all $\sigma \in P$ we have that

$$\frac{|G|}{|G_\sigma|} = \frac{|P|}{|H|}. \tag{2}$$

We claim that

$$\frac{|I \cap G_\sigma|}{|G_\sigma|} = \frac{d_{\sigma^{-1}(q)}}{n} \tag{3}$$

holds for all $\sigma \in P$. Indeed, the previous equality is obvious if $q \notin \{\sigma(1), \ldots, \sigma(m)\}$. Otherwise, let $\sigma(i) = q$. A simple count yields that

$$\frac{|I \cap G_\sigma|}{|G_\sigma|} = \frac{\begin{pmatrix} & & n-1 & & \\ d_1 & d_2 & \ldots & d_i - 1 & \ldots & d_b \end{pmatrix}}{\begin{pmatrix} & & n & & \\ d_1 & d_2 & \ldots & d_i & \ldots & d_b \end{pmatrix}} = \frac{d_i}{n}.$$

We count the elements in I: by using (3) and the fact that for each σ there are exactly $|H|$

permutations ρ such that $G_\rho = G_\sigma$, we get that

$$|I| \cdot |H| = \sum_{\sigma \in P} \frac{d_{\sigma^{-1}(q)}}{n} \cdot |G_\sigma|. \tag{4}$$

It is now easy to combine (2) and (4) to obtain

$$\frac{|I|}{|G|} = \frac{1}{|P|} \sum_{\sigma \in P} \frac{d_{\sigma^{-1}(q)}}{n} \tag{5}$$

whose right-hand term we may rewrite as

$$\frac{1}{|P|} \sum_{\sigma \in P} \frac{d_{\sigma^{-1}(q)}}{n} = \frac{1}{|P|} \sum_{i \in \mathcal{O}} \sum_{\sigma(i)=q} \frac{d_i}{n} = \frac{1}{|P|} \sum_{i \in \mathcal{O}} |\mathrm{Stab}(q)| \frac{d_i}{n} = \frac{1}{|\mathcal{O}|} \sum_{i \in \mathcal{O}} \frac{d_i}{n}. \quad (6) \quad \square$$

Theorem 3.2. *Let $d \le n/2$. If $d_i = d_j$ for all $i, j \in \mathcal{O}$ then the bound in Theorem 3.1 is tight, i.e. the independence ratio of the graph $G = G(P; C; d_1, \ldots, d_b)$ is equal to*

$$\frac{\alpha(G)}{|G|} = \frac{d}{n}.$$

In particular this holds when P is the trivial subgroup of S_b.

Proof. By Theorem 3.1 we have that

$$\frac{\alpha(G)}{|G|} \ge \frac{1}{|\mathcal{O}|} \sum_{i \in \mathcal{O}} \frac{d_i}{n} = \frac{d}{n}.$$

Let L denote the following induced subgraph of G: its vertices are the tuples $(a_1, \ldots, a_n)$ in $V(G)$ that satisfy the following condition: there exists some $0 \le u \le n-1$ such that $a_{i+u} = x_i$ for all $1 \le i \le n$, where

$$(x_1, \ldots, x_n) = (1, \ldots, 1, 2, \ldots, 2, \ldots, m, \ldots, m)$$

is the tuple with d_1 consecutive 1's followed by d_2 consecutive 2's, and so on (sums are understood mod n, as usual). In other words, L consists of all cyclic permutations of the tuple $(x_1, \ldots, x_n)$. In L, let A_u denote the tuple $(a_1, \ldots, a_n)$ such that $a_{i+u} = x_i$ for all $1 \le i \le n$, where the x_i are as above. It is easy to see that the map $u \mapsto A_u$ from $\mathrm{Circ}(d, n)$ to L is a graph isomorphism. It follows from Lemma 2.2 that the stability ratio of L is d/n. Applying Lemma 2.1 to the embedding $L \hookrightarrow G$ we conclude that the stability ratio of G is at most d/n. $\square$

Theorem 3.3. *If $P = S_b$, then the bound in Theorem 3.1 is tight, i.e. the independence ratio of the graph $G = G(P; C; d_1, \ldots, d_b)$ is equal to*

$$\frac{\alpha(G)}{|G|} = \frac{1}{b}.$$

Proof. Clearly $\mathcal{O} = [b] = C$. It follows from Theorem 3.1 that

$$\frac{\alpha(G)}{|G|} \ge \frac{1}{|\mathcal{O}|} \sum_{i \in \mathcal{O}} \frac{d_i}{n} = \frac{1}{b}.$$

Consider the map from the complete graph K_b to G defined by

$$i \mapsto (1+i, \dots, 1+i, 2+i, \dots, 2+i, \dots, m+i, \dots, m+i)$$

where it is understood that $j+i$ appears exactly d_j times and sums are modulo b. It is clear that this is a graph homomorphism and hence it follows from Lemma 2.1 that

$$\frac{1}{b} = \frac{\alpha(K_b)}{|K_b|} \geq \frac{\alpha(G)}{|G|}. \quad \square$$

4. The graph of permutations $G(S_n)$

The next theorem about the structure of stable sets in the graph of permutations appears in the recent paper of Cameron and Ku [4]. However our proof is quite different: we shall deduce it from a more general result on certain subgraphs of $G(S_n)$ (Theorem 4.2). In the next section we shall then generalise this result in various ways, see Theorems 5.1 and 5.7.

It will be convenient in what follows to denote the permutation τ that maps a_i to b_i for $1 \leq i \leq n$ by

$$\tau = \begin{pmatrix} a_1 & a_2 & a_3 & \dots & a_n \\ b_1 & b_2 & b_3 & \dots & b_n \end{pmatrix}.$$

When $a_i = i$ for all $1 \leq i \leq n$, it will also be convenient to denote the permutation τ simply with the n-tuple $(b_1, \dots, b_n)$.

Theorem 4.1. *The sets I_p^q are the only stable sets of maximal size in the graph $G(S_n)$.*

We shall deduce Theorem 4.1 from the more general Theorem 4.2. We extend our definition to the following family of permutation graphs:

Definition. Let $n \geq 2$. Let $1 \leq r \leq n$, and let $1 \leq b_1, \dots, b_r \leq n$ where $b_i \neq b_j$ if $i \neq j$. Define the graph $G(S_n)(b_1, \dots, b_r)$ as the induced subgraph of $G(S_n)$ whose vertices are those permutations σ for which there exists a non-negative integer u, with $0 \leq u \leq n-1$, such that $\sigma(i+u) = b_i$ for all $1 \leq i \leq r$, where sums are understood modulo n, that is the permutations containing the pattern prescribed by $\begin{pmatrix} 1 & 2 & \cdots & r \\ b_1 & b_2 & \cdots & b_r \end{pmatrix}$ or by one of its shifts.

Example 5. $G(S_5)(3, 2, 1)$ consists of all permutations of the form

$(3, 2, 1, x, y)$
$(y, 3, 2, 1, x)$
$(x, y, 3, 2, 1)$
$(1, x, y, 3, 2)$
$(2, 1, x, y, 3)$.

Notice also that if $r = 1$ then $G(S_n)(b_1) = G(S_n)$ for any b_1.

We shall prove the following:

Theorem 4.2. *In $G(S_n)(b_1, \dots, b_r)$ the stable sets of maximal size are of the form $I_p^q \cap G(S_n)(b_1, \dots, b_r)$.*

Lemma 4.3. *Let* $1 \leq s \leq r \leq n$ *and let* $1 \leq j_1 < \cdots < j_s \leq r$*. Then*

(1) $G(S_n)(b_{j_1}, \ldots, b_{j_s}) \supseteq G(S_n)(b_1, \ldots, b_r)$;
(2) $G(S_n)(b_1, \ldots, b_r)$ *contains a copy of* K_n;
(3) $G(S_n)(b_1, \ldots, b_r)$ *is vertex-transitive*;
(4) *the graphs* $G(S_n)(b_1, \ldots, b_r)$ *and* $G(S_n)(1, \ldots, r)$ *are isomorphic, under an isomorphism that preserves the stable sets of the form* $I_p^q \cap G(S_n)(b_1, \ldots, b_r)$.

Proof. The first assertion is clear. Let now consider the subgraph B_u of the graph $G(S_n)(b_1, \ldots, b_r)$ that consists of all permutations σ for which $\sigma(1 + u) = b_1$ (the five tuples listed in the previous Example 5 represent the subgraphs $B_0, \ldots, B_4$ respectively). Clearly the subgraphs B_u's partition $G(S_n)(b_1, \ldots, b_r)$. It is easy to see that $\tau \in B_{u+1}$ if and only if $\tau\gamma \in B_u$, where γ is the n-cycle $(1\,2\,3\, \ldots\, n)$. In particular, $G(S_n)(b_1, \ldots, b_r)$ is closed under right translation by γ and hence contains at least one copy of K_n. If τ is a permutation that fixes every b_i then $\tau\sigma \in B_u$ whenever $\sigma \in B_u$. These permutations act transitively on each block B_u, and it follows that the graph is vertex-transitive. To prove (4), let $\nu \in S_n$ such that $\nu(b_i) = i$ for all $1 \leq i \leq r$, and consider the map $\sigma \mapsto \nu\sigma$: it is easy to verify that it is the required isomorphism. □

Lemma 4.4. *Let* $1 \leq s \leq r \leq n$ *and let* $1 \leq j_1 < \cdots < j_s \leq r$*. If* I *is a stable set of* $G(S_n)(b_{j_1}, \ldots, b_{j_s})$ *of maximal size then* $I \cap G(S_n)(b_1, \ldots, b_r)$ *is a stable set of* $G(S_n)(b_1, \ldots, b_r)$ *of maximal size.*

Proof. We apply Lemma 4.3 to the graphs $G = G(S_n)(b_1, \ldots, b_r)$ and $H = G(S_n)(b_{j_1}, \ldots, b_{j_s})$ to obtain the inclusions

$$K_n \hookrightarrow G \hookrightarrow H \hookrightarrow G(S_n).$$

It follows from Lemmas 2.1 and 4.3 that we have

$$\frac{1}{n} \geq \frac{\alpha(G)}{|V(G)|} \geq \frac{\alpha(H)}{|V(H)|} \geq \frac{\alpha(G(S_n))}{|V(G(S_n))|} = \frac{1}{n}.$$

Hence we have equality and the result follows from the second part of Lemma 2.1 □

The idea of the proof of Theorem 4.2 runs as follows: we use induction on $k = n - r$, the 'degree of freedom' of the graph $G(S_n)(b_1, \ldots, b_r)$. First we prove the case $k = 3$. The second part of the argument is as follows: we take a stable set I of maximal size and intersect it with the subgraphs of smaller 'degree'. We obtain in this way sets of the right shape (i.e. intersections with I_p^q's), and the trick is then to show that we have 'uniformity' among these smaller sets. There are two cases to treat, depending on whether I intersects some subgraph H in a stable set that spreads across several blocks B_u of H or I intersects every subgraph H in a stable set equal to a block of H.

Lemma 4.5. *The statement of Theorem 4.2 holds if* $n - 3 \leq r \leq n$.

Proof. Let $k = n - r$. The result is trivial if $k \leq 2$. Let $k = 3$. Let I be a stable set of maximal size in $G = G(S_n)(b_1, \ldots, b_r)$, which we can take to be, without loss of generality, equal to $G(S_n)(1, \ldots, (n-3))$. Then by the proof of Lemma 4.4 we have that $|I| = 6$. Let $\mu \in B_u$ and $\nu \in B_v$ (notation as in the proof of Lemma 4.3). It is easy to see that, if μ and ν are not adjacent, then $|u - v| \in \{0, 1, 2\} (\mathrm{mod}\, n)$. Hence there exists some u such that I is contained in $B_u \cup \cdots \cup B_{u+4}$, and consequently $|I \cap B_v| \geq 2$ for some v. Let σ and τ denote distinct elements of $I \cap B_v$.

Suppose first that $\sigma(w) \neq \tau(w)$ for all $w \notin \{1 + v, \ldots, (n - 3) + v\}$. We may assume without loss of generality that we have $v = 0$ and the following situation:

$$\begin{array}{rlllllll} \sigma = & (1, & 2, & \ldots, & (n-3), & (n-2), & (n-1), & n) \\ \tau = & (1, & 2, & \ldots, & (n-3), & (n-1), & n, & (n-2)). \end{array}$$

If $\delta \in I \cap B_1$, it is easy to see that

$$\delta = (n, 1, 2, \ldots, (n-3), (n-1), (n-2)).$$

Similarly, we find that $|I \cap B_{n-1}| \leq 1$ and that $|I \cap B_u| = 0$ for $u \notin \{n - 1, 0, 1\}$. Consequently, we must have at least two other vertices in $I \cap B_0$. *In any case*, we will always find vertices σ and τ of $I \cap B_0$ that have exactly $n - 2$ common values.

Case 1. Suppose that the $n - 2$ values where σ and τ coincide are consecutive, i.e. without loss of generality, the following vertices σ and τ are in $I \cap B_0$:

$$\begin{array}{rlllllll} \sigma = & (1, & 2, & \ldots, & (n-3), & (n-2), & (n-1), & n) \\ \tau = & (1, & 2, & \ldots, & (n-3), & (n-2), & n, & (n-1)). \end{array}$$

If $I = B_0$ then $I = I_1^1 \cap G$ and we are done. So suppose that there exists $\mu \in I$, $\mu \notin B_0$. Then $\mu(w) \notin \sigma(w)$ for all $1 \leq w \leq n - 3$, and similarly for τ. It follows easily that $\mu(n - 2) = n - 2$ is forced. If $|I \cap B_0| = 2$ then every permutation in I fixes $n - 2$ and so $I = I_{n-2}^{n-2} \cap G$. Otherwise let $\nu \in I \cap B_0$ distinct from σ and τ. We shall derive a contradiction. We have the following:

$$\begin{array}{rlllllll} \sigma = & (1, & 2, & \ldots, & (n-3), & (n-2), & (n-1), & n) \\ \tau = & (1, & 2, & \ldots, & (n-3), & (n-2), & n, & (n-1)) \\ \nu = & (1, & 2, & \ldots, & (n-3), & x, & y, & z) \end{array}$$

where $x \neq (n - 2)$. Suppose there exists $\mu \in I$ outside B_0; then $\mu(n - 2) = n - 2$, which forces this:

$$\begin{array}{rlllllll} \sigma = & (1, & 2, & \ldots, & (n-3), & (n-2), & (n-1), & n) \\ \tau = & (1, & 2, & \ldots, & (n-3), & (n-2), & n, & (n-1)) \\ \nu = & (1, & 2, & \ldots, & (n-3), & x, & y, & z) \\ \mu = & (2, & \ldots, & (n-3), & x, & (n-2), & y, & 1) \end{array}$$

and we must have $z = n - 2$. Since $|I| = 6$ there must exist another $\nu' \in I \cap B_0$. Then we must have:

$$\begin{array}{rlllllll} \sigma = & (1, & 2, & \ldots, & (n-3), & (n-2), & (n-1), & n) \\ \tau = & (1, & 2, & \ldots, & (n-3), & (n-2), & n, & (n-1)) \\ \nu = & (1, & 2, & \ldots, & (n-3), & x, & y, & (n-2)) \\ \mu = & (2, & \ldots, & (n-3), & x, & (n-2), & y, & 1) \\ \nu' = & (1, & 2, & \ldots, & (n-3), & y, & x, & (n-2)). \end{array}$$

But then μ and ν' are adjacent, a contradiction.

Case 2. Suppose now that the values for which σ and τ coincide are not consecutive; then we have that the following vertices are in I:

$$\begin{array}{lllllllll} \sigma = & (1, & 2, & \ldots, & (n-3), & (n-2), & (n-1), & n) \\ \tau = & (1, & 2, & \ldots, & (n-3), & n, & (n-1), & (n-2)). \end{array}$$

It is easy to verify that every vertex in I outside B_0 must fix $n-1$; it follows that either $I = I_{n-1}^{n-1} \cap G$, or otherwise $I \cap B_0$ contains some other element, whence we are back in Case 1. □

Proof of Theorem 4.2. Let $k = n - r$. We prove the result by induction on k. If $k \leq 3$: this is Lemma 4.5. Now let $k \geq 4$ be an integer such that the result holds for $k-1$. Let I be a stable set of maximal size in $G = G(S_n)(b_1, \ldots, b_r)$, which we may take to be $G(S_n)(1, \ldots, r)$ by Lemma 4.3(4).

Case 1. There exist $\beta \notin \{1, \ldots, r\}$ such that

$$I \cap G(S_n)(1, \ldots, r, \beta) = I_p^q \cap G(S_n)(1, \ldots, r, \beta)$$

for some $q \notin \{1, \ldots, r, \beta\}$.

Let $\sigma \in I$: it is clearly sufficient to prove that $\sigma(p) = q$. There exists some u such that σ looks like this:

$$\sigma = \begin{pmatrix} \ldots & 1+u & 2+u & \ldots & r+u & \ldots \\ \ldots & 1 & 2 & \ldots & r & \ldots \end{pmatrix}.$$

We build a permutation $\tau \in I$ as follows: (i) let $\tau(p) = q$. (ii) Next, notice that since $n - r \geq 4$, there exist at least 3 distinct elements $v \in \{1, 2, \ldots, n\}$ such that $p \notin \{1+v, \ldots, r+v, (r+1)+v\}$. Choose one of these v such that $\sigma(1+v) \neq 1$ and $\sigma((r+1)+c) \neq \beta$. Now define $\tau(i+v) = i$ for $1 \leq i \leq r$ and $\tau((r+1)+v) = \beta$. (iii) There are at least 2 elements in $\{1, 2, \ldots, n\}$ where τ is not yet defined: choose its values so that it has no value in common with σ. Now it is clear by construction that $\tau \in I_p^q \cap G(S_n)(1, \ldots, r, \beta)$ and hence $\tau \in I$, and that $\sigma(x) \neq \tau(x)$ for all $x \neq p$. Since σ and τ are not adjacent it follows that $\sigma(p) = q$.

By induction hypothesis and Lemma 4.4, we are left with this:

Case 2. For every $\beta \notin \{1, \ldots, r\}$ there exists p and there exists $q \in \{1, \ldots, r, \beta\}$ such that

$$I \cap G(S_n)(1, \ldots, r, \beta) = I_p^q \cap G(S_n)(1, \ldots, r, \beta).$$

By permuting and renaming entries we may assume without loss of generality that the identity permutation

$$\alpha = (1, 2, \ldots, r, r+1, \ldots, n) \in I.$$

Then α belongs to

$$I \cap G(S_n)(1, \ldots, r, r+1) = I_p^q \cap G(S_n)(1, \ldots, r, r+1), \tag{7}$$

and it is easy to see that in fact we may take $p = q = 1$. In particular, I contains all permutations that fix every $1 \leq i \leq r+1$. We shall show that

$$I = I_1^1 \cap G(S_n)(1, \ldots, r).$$

Suppose for a contradiction that there exists $\sigma \in I$ such that $\sigma(1) \neq 1$; so there exists $u \neq 0$ such that $\sigma(i+u) = i$ for all $1 \leq i \leq r$. Notice that (7) implies that $\sigma((r+1)+u) = \beta \neq r+1$. Now

$$\sigma \in I \cap G(S_n)(1,2,\ldots,r,\beta) = I_{p'}^{q'} \cap G(S_n)(1,2,\ldots,r,\beta)$$

for some $q' \in \{1,\ldots,r,\beta\}$ so as above we conclude that I contains every permutation τ such that $\tau(i+u) = i$ for $i = 1,\ldots,r$ and $\tau((r+1)+u) = \beta$. Since there are at least 2 entries not in $\{1+u,\ldots,(r+1)+u\}$, we can find such a permutation τ which is fixed-point free, contradicting the fact that $\alpha \in I$, unless $\beta = (r+1)+u$; but then $\beta \notin \{1,2,\ldots,r,r+1\}$, so we may choose a permutation $\mu \in I$ that fixes $1 \leq i \leq r+1$ such that $\tau(x) \neq \mu(x)$ for all $x \in [n]$ and we are done. □

5. The case $P = S_b$

Now we investigate the shape of the maximal stable sets in the case where $P = S_b$: notice that in this case we must have $C = [b]$. First we restrict our attention to the case when d_i is equal to 0 or 1 for all i, with $b \geq n$. More precisely, recall from Example 4 that $G(S_b)^{(n)}$ denotes the graph $G(S_b;[b];1,1,\ldots,1,0,\ldots,0)$ (where 1 appears n times). The vertices of this graph are the injections of $[n]$ into $[b]$. We shall require in this section the following generalisation of Theorem 4.1:

Theorem 5.1. *The sets I_p^q are the only stable sets of maximal size in $G(S_b)^{(n)}$.*

Proof. This is an easy application of the no-homomorphism lemma. By Theorem 3.3 the stability ratio of $G = G(S_b)^{(n)}$ is equal to $1/b$. It is obvious that the following map is a surjective graph homomorphism from $G(S_b)$ onto $G(S_b)^{(n)}$:

$$\phi : \sigma \longmapsto (\sigma(1),\ldots,\sigma(n)).$$

If $b = 2$ the result is trivial, so assume from now on that $b \geq 3$. Let I be a stable set of maximal size in $G(S_b)^{(n)}$. By the second part of Lemma 2.1, we have that $\phi^{-1}(I)$ is a stable set of maximal size in $G(S_b)$. By Theorem 4.1 it must be of the form I_p^q. Since $b \geq 3$ we can find $\sigma, \tau \in I_p^q$ such that $\sigma(x) \neq \tau(x)$ for all $x \neq p$. Since $\phi(\sigma)$ is not adjacent to $\phi(\tau)$, it follows that $p \leq n$. Since ϕ is onto it follows that $a_p = q$ for all $(a_1,\ldots,a_n) \in I$, and thus $I = I_p^q$ in $G(S_b)^{(n)}$. □

Now we investigate the shape of the maximal stable sets in the case of a fixed graph $G = G(P;C;d_1,\ldots,d_b)$ still in the case where P is the full symmetric group on b letters, but for a more general sequence of the frequencies d_i; as mentioned earlier we have that $C = [b]$.

- Notice that if $d_1 > n/2$ then there are stable sets of maximal size of the 'wrong' form: indeed, take I to be the set that consists of all tuples that contain the entry 1 d_1 times.
- If the least non-zero d_i is equal to $n/2$, then of course we have $d_1 = d_2 = n/2$ and $d_i = 0$ for all $i \geq 3$. Then there are also stable sets of maximal size of the 'wrong' shape: consider the set I of all tuples that have $n/2$ 1's appearing in the first

$n-1$ coordinates. Clearly this is a stable set, and it does not have the right form. A simple count yields that it has maximal cardinality.

Hence, from now on, we assume that each d_i is at most $n/2$ and that the number of non-zero d_i is at least 3. We shall also assume that $b \geq 5$ (see concluding remarks).

Let I be a stable set in G of maximal size. Recall that m is the largest index such that d_m is non-zero. We shall proceed as follows: fix a 'pattern', i.e. a partition θ of the index set $[n]$, with blocks of the required size $d_1, d_2, \ldots, d_m$, and let G_θ denote the subgraph of G that consists of all tuples $(a_1, \ldots, a_n)$ such that $a_i = a_j$ if and only if i and j lie in the same block of θ. It is easy to see that this graph is isomorphic to $G(S_b)^{(m)}$: indeed, just reorder the indices to get tuples of the form $(\sigma(1), \ldots, \sigma(1), \ldots, \sigma(m), \ldots, \sigma(m))$ and then identify equal coordinates. It follows from Lemma 2.1 that the inverse image of I under this isomorphism is a maximal stable set of $G(S_b)^{(m)}$, and we know these have the right form by Theorem 5.1. In other words, for every partition θ, there exists an index p and a value q such that $G_\theta \cap I$ consists of all tuples whose pth coordinate is equal to q. As usual, we shall denote this set by $I_p^q \cap G_\theta$. We shall prove that the values of p and q are independent of θ. Here is the strategy. Let θ and θ' be partitions of $[n]$ with m blocks $B_1, \ldots, B_m$ and $B'_1, \ldots, B'_m$ respectively. We shall say that these partitions *differ by a transposition* if there exist distinct elements i and j of $[n]$ and distinct indices u, v such that $B'_u = B_u \cup \{i\} \backslash \{j\}$, $B'_v = B_v \cup \{j\} \backslash \{i\}$ and $B_k = B'_k$ for all k distinct from u and v (in other words, we obtain the partition θ' from θ by choosing two blocks of θ and exchanging two elements, one from each). It is clear that for every pair of partitions θ and θ', there exists a sequence of partitions

$$\theta = \theta_0, \ldots, \theta_l = \theta'$$

such that θ_i and θ_{i+1} differ by a transposition for every i.

One way of seeing this is as follows: given a partition θ, consider the n-tuple X_θ whose ith coordinate is the symbol B_j if i is in block B_j of the partition. Then we may transform the tuple for θ into the tuple for θ' simply by a series of transpositions of two symbols. Notice that with this notation, G_θ is simply the set of all n-tuples obtained from X_θ by assigning distinct values to the symbols $B_1, \ldots, B_m$.

We shall first show the following:

Lemma 5.2. *Let $b \geq 4$. If θ and θ' are partitions that differ by a transposition, and $I \cap G_\theta = I_p^q \cap G_\theta$ and $I \cap G_{\theta'} = I_{p'}^{q'} \cap G_\theta$, then $q = q'$.*

Proof. Suppose by contradiction that $q \neq q'$. Without loss of generality suppose that $q = 1$ and $q' = 2$. A case-by-case analysis of the relative positions of p and p' shall exhibit elements of $G_\theta \cap I$ and $G_{\theta'} \cap I$ that are adjacent. We use the tuples X_θ and $X_{\theta'}$ for ease of discussion. Without loss of generality, suppose that the blocks in which θ and θ' differ are B_1 and B_2, so that tuples X_θ and $X_{\theta'}$ look something like this:

$$\begin{matrix} (B_1, & B_1, & \ldots, & B_1, & B_1, & B_2, & B_2, & \ldots, & B_l) \\ (B_1, & B_1, & \ldots, & B_1, & B_2, & B_1, & B_2, & \ldots, & B_l) \end{matrix}.$$

Case 1. Suppose that neither p nor p' are in blocks B_1, B_2.

Subcase 1. Suppose that p and p' are in the same block (of θ and/or θ'), say B_3. Then the following assignment of values leads to a contradiction:

$$\begin{array}{l} (2, \ \ldots, \ \ldots, \ 2, \ 3, \ \ldots, \ \ldots, \ 3, \ 1, \ \ldots, \ 1, \ B_4, \ \ldots) \\ (4, \ \ldots, \ 4, \ 1, \ 4, \ 1, \ \ldots, \ 1, \ 2, \ \ldots, \ 2, \ B_4, \ \ldots) \end{array}$$

where the symbols $B_4, B_5, \ldots$ are filled as follows: if there is only one more block, put the value 4 for X_θ and the value 3 for $X_{\theta'}$; if there is more than one block, simply assign any values to $B_4, B_5, \ldots$ of X_θ and cycle them to give an assignment to $B'_4, \ldots$ of $X_{\theta'}$ so that no two coordinates are equal.

Subcase 2. Suppose that p and p' are in different blocks, say B_3 and B_4. Then the following assignment does the job:

$$\begin{array}{l} (2, \ \ldots, \ \ldots, \ 2, \ 3, \ \ldots, \ \ldots, \ 3, \ 1, \ \ldots, \ 1, \ 4, \ \ldots, \ 4, \ B_5, \ \ldots) \\ (4, \ \ldots, \ 4, \ 1, \ 4, \ 1, \ \ldots, \ 1, \ 3, \ \ldots, \ 3, \ 2, \ \ldots, \ 2, \ B_5, \ \ldots) \end{array}$$

where the blocks $B_5, \ldots$ are filled as in Subcase 1. (Note that if $b = 5$ and there are exactly 5 blocks, it is necessary to switch some values, say setting B_4 to 5 instead of 4 in the first tuple, etc. but this is easy.)

Case 2. Suppose that p is in B_1 but p' is not in blocks B_1, B_2. Assume that p' is in block B_3. Then the following tuples do the job:

$$\begin{array}{l} (1, \ \ldots, \ \ldots, \ 1, \ 2, \ \ldots, \ \ldots, \ 2, \ 3, \ \ldots, \ 3, \ B_4, \ \ldots) \\ (4, \ \ldots, \ 4, \ 3, \ 4, \ 3, \ \ldots, \ 3, \ 2, \ \ldots, \ 2, \ B_4, \ \ldots) \end{array}$$

where the blocks $B_4, \ldots$ are filled as before.

Case 3. By symmetry, it remains to check this case: both p and p' are in B_1 or B_2.

Subcase 1. Both p and p' are in B_1. This is easy, consider the tuples

$$\begin{array}{l} (1, \ \ldots, \ \ldots, \ 1, \ 4, \ \ldots, \ \ldots, \ 4, \ B_3, \ \ldots) \\ (2, \ \ldots, \ 2, \ 3, \ 2, \ 3, \ \ldots, \ 3, \ B_3, \ \ldots) \end{array}$$

where the blocks $B_3, \ldots$ are easily filled.

Subcase 2. p is in B_1 and p' is in B_2. This is similar to Subcase 1. □

Remark. It follows from the lemma and the preceding remarks that we may assume without loss of generality that $q = 1$ for all partitions.

Lemma 5.3. *Let $b \geq 5$. If θ and θ' are partitions that differ by a transposition, and $I \cap G_\theta = I_p^1 \cap G_\theta$ and $I \cap G_{\theta'} = I_{p'}^1 \cap G_\theta$, then the block of θ that contains p and the block of θ' that contains p' must intersect.*

Proof. We follow a similar case-by-case procedure. Suppose by contradiction that the blocks do not intersect.

Case 1. Suppose that p and p' are neither in B_1 nor B_2 (as before we assume that the partitions differ only in these blocks). So assume without loss of generality that p is in B_3

and p' is in B_4. Then the following tuples are adjacent:

$$\begin{array}{cccccccccccccccc} (3, & \ldots, & \ldots, & 3, & 4, & \ldots, & \ldots, & 4, & 1, & \ldots, & 1, & 2, & \ldots, & 2, & B_5, & \ldots) \\ (2, & \ldots, & 2, & 5, & 2, & 5, & \ldots, & 5, & 3, & \ldots, & 3, & 1, & \ldots, & 1, & B_5, & \ldots) \end{array}$$

where the blocks $B_5, \ldots$ are filled as before.

Case 2. Suppose that p is in B_1 and p' is in B_3. Then the following tuples do the job:

$$\begin{array}{cccccccccccccc} (1, & \ldots, & \ldots, & 1, & 4, & \ldots, & \ldots, & 4, & 2, & \ldots, & 2, & B_4, & \ldots) \\ (2, & \ldots, & 2, & 3, & 2, & 3, & \ldots, & 3, & 1, & \ldots, & 1, & B_4, & \ldots) \end{array}$$

where the blocks $B_4, \ldots$ are filled as before.

Case 3. By symmetry, it would remain to check this case: p is in B_1 and p' is in B_2. But the block B_1 of θ and the block B_2 of θ' share a coordinate, so we are done. □

The preceding lemma has the following consequence: let θ be a partition such that $I \cap G_\theta = I_p^1 \cap G_\theta$, and let B denote the block of θ that contains p. Let θ' be a partition obtained from θ by a transposition *outside* the block B, i.e. θ and θ' differ by a transposition *and* both contain the block B. Let $I \cap G_{\theta'} = I_{p'}^1 \cap G_{\theta'}$. By the last lemma, the block of θ' that contains p' must intersect block B, and hence it must be equal to B. Now, any partition α that contains the block B can be obtained from θ via a sequence of transpositions; hence we must have $I \cap G_\alpha = I_p^1 \cap G_\alpha$. So we have proved this:

Lemma 5.4. *Let θ be a partition such that $I \cap G_\theta = I_p^1 \cap G_\theta$, and let B denote the block of θ that contains p. Let α be any other partition (with correct block sizes) which contains block B. Then $I \cap G_\alpha = I_p^1 \cap G_\alpha$.* □

Lemma 5.5. *Let $b \geq 5$. If θ and θ' are any partitions with $I \cap G_\theta = I_p^1 \cap G_\theta$ and $I \cap G_{\theta'} = I_{p'}^1 \cap G_\theta$, then the block of θ that contains p and the block of θ' that contains p' must intersect.*

Proof. Suppose for a contradiction that this is not the case. Let θ and θ' be partitions that witness this: $I \cap G_\theta = I_p^1 \cap G_\theta$ and $I \cap G_{\theta'} = I_{p'}^1 \cap G_\theta$, let B denote the block of θ that contains p and let B' denote the block of θ' that contains p'; suppose then that B and B' are disjoint.

Case 1. Suppose that the blocks B and B' have different cardinalities. Then we may certainly find a partition α which contains both these blocks and has the correct block sizes. It is clear that the existence of α contradicts Lemma 5.4.

Case 2. Suppose now that the blocks B and B' have the same cardinality. We will construct n-tuples $X = (x_1, \ldots, x_n)$ and $Y = (y_1, \ldots, y_n)$, both in I but adjacent. Let X be any tuple (with the correct block sizes) such that $x_i = 1$ for all $i \in B$. Consider the tuple Z obtained from X by 'swapping' the entries in blocks B and B' (the order of the entries is immaterial). Now define Y simply by replacing every entry z of Z by $z+1$ if $2 \leq z \leq b-1$ and by 2 if $z = b$. Clearly the tuples X and Y are adjacent. However, the partition α such that $X \in G_\alpha$ contains the block B, hence by Lemma 5.4 we have that $X \in I$. The same argument using block B' shows that $Y \in I$, a contradiction. □

Lemma 5.6. *Let $m \geq 3$, let $d_1 \geq d_2 \geq \cdots \geq d_m$ be positive integers such that*

$$\sum_{i=1}^{m} d_i = n \geq 1,$$

$$d_m < n/2$$

and

$$d_1 \leq n/2.$$

Let F be a family of intersecting subsets of $[n]$ such that both the following conditions hold:

(i) *each set in F is a block of some partition θ of $[n]$ with blocks of size $d_1, \ldots, d_m$*
(ii) *each such partition θ has one of its blocks in F.*

Then there exists an $i \in [n]$ such that F consists of all subsets of $[n]$ of one of the prescribed sizes $d_1, \ldots, d_m$ that contain i.

Proof. Let $P(n; d_1, \ldots, d_m)$ denote the number of partitions of $[n]$ into blocks of sizes $d_1, \ldots, d_m$. Let $\{d_{i_1} > \cdots > d_{i_k}\}$ be the *distinct* values that the d_i's take. We have a recurrence:

$$P(n; d_1, \ldots, d_m) = \sum_{j=1}^{k} \binom{n-1}{d_{i_j}-1} P(n - d_{i_j}; d_1, \ldots, d_{i_j-1}, d_{i_j+1}, \ldots, d_m). \qquad (8)$$

This is easy to obtain: to each partition θ with the right block sizes, we associate its block (say of size d_{i_j}) that contains the element 1; the rest of the partition is a partition of the remaining $n - d_{i_j}$ elements in blocks of the remaining sizes.

Consider the map that assigns to every partition θ of $[n]$ with blocks of size $d_1, \ldots, d_m$ the (obviously unique) block f_θ of θ which is in F. Let $f \in F$. Of course if $f = f_\theta$ then f is a block of θ. Conversely, if f happens to be a block of some partition θ then we must have $f_\theta = f$ since every pair of members of F intersect. Hence the partitions that map to f are precisely those that have f as a block: if f has size d_i then there are $P(n - d_i; d_1, d_2, \ldots, d_{i-1}, d_{i+1}, \ldots, d_m)$ of these.

Let F_j denote the set of members of F with cardinality d_{i_j}, and let α_j denote the number of elements in F_j. Then the above argument shows that

$$P(n; d_1, \ldots, d_m) = \sum_{j=1}^{k} \alpha_j P(n - d_{i_j}; d_1, d_2, \ldots, d_{i_j-1}, d_{i_j+1}, \ldots, d_{i_k}). \qquad (9)$$

Since F_j is an intersecting family, the Erdős–Ko–Rado inequality [7] tells us that

$$\alpha_j \leq \binom{n-1}{d_{i_j}-1}$$

for every j. It follows from Eqs. (8) and (9) that

$$\alpha_j = \binom{n-1}{d_{i_j}-1}$$

for every j. Fix $j_0 < n/2$. By the Hilton–Milner inequality [10] we have that F_{j_0} consists of all sets that contain some fixed value s. Let $X \in F$. Since $|X| + j_0 \le n$, it is clear that, if X intersects every set in F_{j_0}, then it must contain s. Hence for every j, F_j consists only of sets that contain the element s, and the computation of α_j above shows that in fact every set of cardinality d_{i_j} that contains s must be in F_j. □

Theorem 5.7. *Let $b \ge 5$, and let $n/2 \ge d_1 \ge d_2 \ge \cdots d_b \ge 0$ be integers, with $n = \sum_i d_i \ge 1$. If there are at least 3 non-zero d_i's, then the sets I_p^q are the only stable sets of maximal size in the graph $G(S_b; C; d_1, \ldots, d_b)$.*

Proof. Let I be a stable set of maximal size in $G(S_b; C; d_1, \ldots, d_b)$. By Lemma 5.2 there exists a unique q such that $I \cap G_\theta = I^q_{p(\theta)} \cap G_\theta$ for all partitions θ. For each θ, let B_θ denote the block of θ which contains $p(\theta)$, and let F be the set that consists of all these blocks. By Lemma 5.5 F satisfies the hypotheses of Lemma 5.6, and thus there exists a unique p such that these blocks are precisely those that contain p. □

6. Concluding remarks

The case $P = S_b$ with $b \le 4$

Although we have verified a few small cases when $b \in \{3, 4\}$ it is apparent that the method we used for $b \ge 5$ fails in these cases and that a different approach is required. Notice that the cases where $P = S_b$ or P is trivial are related, as for example in the case where $d_1 = d_2 = \cdots = d_b$ and we obtain the same graph, provided $C = [b]$.

The case $P = 1$

When the group P is trivial, we have the following result:

Theorem 6.1 ([12]). *If P is trivial and $d_1 > d_2$ the I_p^1 are the only stable sets of maximal size in the graph $G(P; [b]; d_1, \ldots, d_b)$.*

From this one can easily deduce (yet another) characterisation of the stable sets of maximal size in Kneser graphs: let $1 < r < n/2$ and consider the graph $G(1; [b]; r, 1, \ldots, 1)$ where 1 appears $n - r$ times, and the homomorphism

$$\phi : G(1; [b]; r, 1, \ldots, 1) \to K(r, n)$$

which maps a tuple $(x_1, \ldots, x_n)$ to the set of indices i such that $x_i = 1$. Let I be a stable set of maximal size in $K(r, n)$: since these graphs have the same independence ratio, the no-homomorphism lemma and the last result guarantee that $\phi^{-1}(I) = I_p^1$, which means that all the members of I contain p. Recent results on stable sets of maximal size in powers of regular graphs, including the case of Kneser graphs, may be found in [1].

Automorphisms of G

It is easy to verify that under mild conditions both the actions of S_n and P are faithful on the graph $G = G(P; C; d_1, \ldots, d_b)$: all that is required is that $m \ge 2$ and that the union of the orbits of $1 \le i \le m$ under P is equal to $[b]$ (which can always be assumed).

Furthermore if some d_i is not equal to 1 then in fact the group of automorphisms of G will contain a copy of $S_n \times P$. If the I_p^q are the only stable sets of maximal size in G, then under what conditions is $\mathrm{Aut}(G) = S_n \times P$?

Chromatic number of G

If P is trivial, then the special case of Kneser graphs shows that the chromatic number of G may be quite difficult to compute [11, 14]. On the other hand, when $P = S_b$ the graph contains a copy of K_b, and the projection on a coordinate shows that its chromatic number is b. What can be said for more general subgroups P of S_b?

Acknowledgements

The authors wish to thank János Körner for suggesting this problem and Claude Tardif and Daniele A. Gewurz for helpful comments. This research was completed during the first author's stays at the Dipartimento di Informatica, Università di Roma "La Sapienza" in 2001 and 2003. The first author's research is supported by a grant from NSERC.

References

[1] N. Alon, I. Dinur, E. Friedgut, B. Sudakov, Graph Products, Fourier Analysis and Spectral Techniques, 2003, p. 22 (preprint).
[2] M.O. Albertson, K.L. Collins, Homomorphisms of 3-chromatic graphs, Discrete Math. 54 (1985) 127–132.
[3] B. Bollobás, Combinatorics, Set Systems, Hypergraphs, Hamilies of Vectors and Combinatorial Probability, Cambridge University Press, 1986.
[4] P.J. Cameron, C.Y. Ku, Intersecting families of permutations, European J. Combin. 24 (2003) 881–890.
[5] M. Deza, P. Frankl, On the maximum number of permutations with given maximal or minimal distance, J. Combin. Theory Ser. A 22 (1977) 352–360.
[6] M. Deza, P. Frankl, Problem, in: Combinatorics, Coll. Math. Soc. János. Bolyai, North-Holland, Amsterdam, 18 (1978) p. 1193.
[7] P. Erdős, C. Ko, R. Rado, Intersection theorems for systems of finite sets, Quart. J. Math. Oxford Ser. (2) 12 (1961) 313–318.
[8] C. Godsil, G. Royle, Algebraic Graph Theory, Graduate Texts in Mathematics, Springer, New York, Berlin, Heidelberg, 2001.
[9] D. Greenwell, L. Lovász, Applications of product colouring, Acta Math. Acad. Sci. Hungar. 25 (1974) 335–340.
[10] A.J.W. Hilton, E.C. Milner, Some intersection theorems for systems of finite sets, Quart. J. Math. Oxford Ser. (2) 18 (1967) 369–384.
[11] M. Kneser, Aufgabe 360, Jahresbes Deutsch. Math.—Verein 58(2 Abt.) p. 27 (1955).
[12] B. Larose, Some results on stable sets of maximal size in vertex-transitive graphs (in preparation).
[13] B. Larose, C. Tardif, Projectivity and independent sets in powers of graphs, J. Graph Theory 40 (3) (2002) 162–171.
[14] L. Lovász, Kneser's conjecture, chromatic number and homotopy, J. Combin. Theory Ser. (A) 25 (1978) 319–324.