Available online at www.sciencedirect.com
SCIENCE DIRECT®

Journal of Combinatorial Theory, Series B 92 (2004) 359–377

Journal of Combinatorial Theory Series B

www.elsevier.com/locate/jctb

Specht modules and chromatic polynomials

Norman Biggs

Centre for Discrete and Applicable Mathematics, London School of Economics, Houghton Street, London WC2A 2AE, UK

Received 17 November 2003

Abstract

An explicit formula for the chromatic polynomials of certain families of graphs, called 'bracelets', is obtained. The terms correspond to irreducible representations of symmetric groups. The theory is developed using the standard bases for the Specht modules of representation theory, and leads to an effective means of calculation.

MSC: 05C15; 05C50

Keywords: Chromatic polynomial; Standard tablean

1. Introduction

The *chromatic polynomial* $P(G;k)$ is the function which gives the number of ways of colouring a graph G when k colours are available. The fact that it is a polynomial function of k is elementary (Section 2), related to the fact that, when k is large enough, not all the colours can be used. Another quite trivial property of the construction is that the names of the k colours are immaterial; in other words, if we are given a colouring, then any permutation of the colours produces another colouring. In Section 2, these facts will be cast into an algebraic form that provides the foundation of our theory.

A 'bracelet' $G_n = G_n(B, L)$ is formed by taking n copies of a graph B and joining each copy to the next by a set of links L (with $n + 1 = 1$ by convention). Using the framework described in Section 2, it can be shown that the chromatic polynomial of G_n

E-mail address: n.l.biggs@lse.ac.uk.

0095-8956/$ - see front matter
doi:10.1016/j.jctb.2004.09.001

can be expressed in the form

$$P(G_n; k) = \sum_{\pi} m_{B,\pi}(k)\,\mathrm{tr}(N_L^{\pi})^n.$$

The sum is taken over all partitions π such that $0 \leqslant |\pi| \leqslant b$, where b is the number of vertices of B. The terms $m_{B,\pi}(k)$ are polynomials in k, and they are independent of L. When B is the complete graph K_b the relevant polynomials $m_{\pi}(k)$ are given by a remarkably simple formula (see Sections 3 and 5).

The size of the matrix N_L^{π} is independent of k; its entries are polynomials in k, and they do depend on L. The original approach to these matrices [3] involved a sequence of elementary, but complicated, calculations, culminating in a rather mysterious application of representation theory. Here we shall present the theory in a more elegant form. In Sections 3 and 4, we construct bases for certain irreducible modules (corresponding to the Specht modules of representation theory), and we shall use these bases for our calculations.

The results obtained here also facilitate further study of the general properties of the matrices N_L^{π}. In particular, we are strongly motivated by the fact that the formula displayed above is well adapted to the application of the Beraha–Kahane–Weiss theorem [1], leading to the construction of 'equimodular curves' [4] that describe the behaviour of the roots of $P(G_n; k)$ for large values of n.

2. Colourings and modules

Let B be a graph with vertex-set V and edge-set E. A *colour-partition* of B is a partition of V into independent sets:

$$\mathcal{P} = \{P_1, P_2, \ldots, P_r\}.$$

A k-colouring of B is a function $c : V \to K$, where $K = \{1, 2, \ldots, k\}$, such that $c(v) \neq c(w)$ whenever $vw \in E$. Clearly, any k-colouring induces a colour-partition, each part being a set of vertices that are assigned a particular colour. A colour-partition with $|\mathcal{P}|$ parts is induced by

$$(k)_{|\mathcal{P}|} = k(k-1)\ldots(k-|\mathcal{P}|+1),$$

k-colourings, so the total number of k-colourings is

$$P(B; k) = \sum_{\mathcal{P}} (k)_{|\mathcal{P}|} = \sum_{r=1}^{|V|} q_r(B)\ (k)_r,$$

where $q_r(B)$ the number of colour-partitions of B with r parts. This simple argument shows that $P(B; k)$ is a polynomial function of k. For our purposes we require its algebraic counterpart, as follows.

Denote by $\mathcal{V}_k(B)$ the complex vector space with basis the set of all k-colourings of B. Clearly, it is the direct sum of subspaces

$$\mathcal{V}_k(B) = \bigoplus \mathcal{V}_{k,\mathcal{P}},$$

where $\mathcal{V}_{k,\mathcal{P}}$ is the subspace whose basis is the set of k-colourings that induce $\mathcal{P}$. The symmetric group Sym_k of all permutations of the set $\{1, 2, \ldots, k\}$ acts on $\mathcal{V}_k(B)$ by the rule $\omega(c) = \omega c$, which makes $\mathcal{V}_k(B)$ a $\mathbb{C}\mathrm{Sym}_k$-module (For the avoidance of doubt, we state that, in this paper, the composite of two permutations ω_1, ω_2 is given by $(\omega_1\omega_2)(x) = \omega_1(\omega_2(x))$.) This action preserves the subspaces $\mathcal{V}_{k,\mathcal{P}}$, and so they are $\mathbb{C}\mathrm{Sym}_k$-submodules.

Of course, $\mathcal{V}_{k,\mathcal{P}}$ is just the module generated by the injections of an r-set into a k-set, and its decomposition is an exercise in the representation theory of the symmetric group [8,11]. The analysis will be done here in terms that allow us to appeal directly to the results as they are stated in [11], although we shall introduce some minor modifications to the terminology.

A partition λ of a nonnegative integer k is a sequence $(\lambda_1, \lambda_2, \ldots, \lambda_k)$ such that

$$\lambda_1 + \lambda_2 + \cdots + \lambda_k = k, \quad (\lambda_1 \geqslant \lambda_2 \geqslant \cdots \lambda_k \geqslant 0).$$

The notation is often abbreviated by collecting equal parts and omitting the parts that are zero: for example $(4^2, 3)$ is a partition of 11 with three non-zero parts 4,4, and 3. Associated with λ is a *diagram* composed of *cells* (i, j) arranged in rows and columns: there are λ_i cells $(i, 1), (i, 2), \ldots, (i, \lambda_i)$ in row i (see below for examples). We denote the set of cells by $[\lambda]$. Conventionally, there are no cells corresponding to parts of λ that are zero; in particular when $k = 0$ we have the partition o for which $[o] = \emptyset$.

Given a partition λ we define a λ-*tableau* to be a function $t : [\lambda] \to \mathbb{N} \cup \{0\}$. Note that this corresponds to Sagan's [11, 2.9.1] 'generalized Young tableau' except that we allow the value 0 as well as positive integers. A tableau is represented by putting the values in the appropriate cells: for example, if $\lambda = (4^2, 3)$, the following is a λ-tableau:

$$\begin{array}{cccc} 0 & 2 & 5 & 3 \\ 7 & 3 & 2 & 0 \\ 1 & 3 & 6 & \end{array}.$$

The link with graph colourings depends on the simple observation that a k-colouring c of a graph B, which induces a colour partition $\mathcal{P}$ with $r = |\mathcal{P}|$, can be represented (provided $k > r$) by a tableau corresponding to the partition $\lambda_{k,r} = (k - r, 1^r)$:

$$\begin{array}{ccccc} * & * & * & \cdots & * \\ * & & & & \\ \cdot & & & & \\ \cdot & & & & \\ \cdot & & & & \\ * & & & & \end{array}.$$

Here each $*$ stands for one of the colours, that is, the numbers $1, 2, \ldots, k$. The $k - r$ colours in the top row are those that c does not assign to any vertex. There is one colour in each of the remaining rows, these colours being the ones that c assigns to the independent sets comprising $\mathcal{P}$. Note that this is a *bijective* tableau on $\{1, 2, \ldots, k\}$; in other words, each value occurs exactly once in a cell.

In order to take this idea further, we need some more terminology. We shall denote the rows of $[\lambda]$ by r_i $(i = 0, 1, 2, \ldots)$, and the columns by c_j $(j = 1, 2, \ldots)$. Thus

$$[\lambda] \;=\; r_0 \cup r_1 \cup r_2 \cup \ldots = c_1 \cup c_2 \cup \ldots \quad .$$

The reason for calling the top row r_0 will appear later. The *row stabilizer* and *column stabilizer* corresponding to λ are defined to be, respectively, the subgroups R_λ and C_λ of the symmetric group $\mathrm{Sym}[\lambda]$ of permutations of $[\lambda]$, given by

$$R_\lambda = \mathrm{Sym}(r_0) \times \mathrm{Sym}(r_1) \times \ldots \qquad \text{and} \qquad C_\lambda = \mathrm{Sym}(c_1) \times \mathrm{Sym}(c_2) \times \ldots .$$

Given a λ-tableau t and $\rho \in R_\lambda$, $t\rho$ is a λ-tableau in which the values occurring in each row are the same as those in t, but in a different order. In the case when t is a bijective λ-tableau on $\{1, 2, \ldots, k\}$, the equivalence class

$$\{t\} = \{t\rho \mid \rho \in R_\lambda\}$$

is known as a *tabloid* [11, 2.1.4].

Let $\mathcal{Z}^\lambda$ denote the complex vector space with basis the set of all bijective λ-tableaux on $\{1, 2, \ldots, k\}$. Associated with each tabloid we have an element of $\mathcal{Z}^\lambda$:

$$\{t\} \quad \longleftrightarrow \quad f_t = \sum_{s \in \{t\}} s = \sum_{\rho \in R_\lambda} t\rho.$$

The space spanned by these elements will be denoted by $\mathcal{M}^\lambda$ (In the usual development of the subject [11, 2.1.5] $\mathcal{M}^\lambda$ is defined directly as the complex vector space with basis the set of tabloids.) Note that $\mathcal{M}^\lambda$ is a $\mathbb{C}\,\mathrm{Sym}_k$-module by virtue of the action of Sym_k on $\{1, 2, \ldots, k\}$.

In the correspondence between colourings and tableaux described above, it is clear that order of the numbers within each row is irrelevant. So each k-colouring c corresponds to a $\lambda_{k,r}$-tabloid, where $r = |\mathcal{P}|$ is the number of colours actually used in c. We have the isomorphism

$$\mathcal{V}_{k,\mathcal{P}} \approx \mathcal{M}^{\lambda_{k,r}}.$$

It is a standard result [11, 2.4.7] that, for any partition λ of k, the irreducible constituents of the $\mathbb{C}\mathrm{Sym}_k$-module $\mathcal{M}^\lambda$ are *Specht modules* S^μ, where μ is a partition that *dominates* λ. This means that

$$\mu_1 + \mu_2 + \cdots + \mu_i \;\geqslant\; \lambda_1 + \lambda_2 + \cdots + \lambda_i \quad (i = 1, 2, \ldots k).$$

When $\lambda = \lambda_{k,r}$, the condition with $i = 1$ implies that $\mu_1 \geqslant k - r$. Writing $\mu_1 = k - \ell$, $(0 \leqslant \ell \leqslant r)$, it follows that the remaining conditions are satisfied when $\pi = (\mu_2, \mu_3, \ldots)$ is any partition of ℓ. Thus, provided k is large enough, the partitions μ of k that dominate $\lambda_{k,r}$ are in bijective correspondence with the partitions π such that $0 \leqslant |\pi| \leqslant r$. The inverse bijection is such that, given π such that $|\pi| = \ell$, the corresponding partition of k is

$$\pi^k = (k - \ell, \pi_1, \pi_2, \ldots, \pi_\ell) \quad (k \geqslant 2\ell).$$

With this notation, the foregoing results can be summarized as follows.

Lemma 1. *For all $k \geqslant 2|\mathcal{P}|$, $\mathcal{V}_{k,\mathcal{P}}$ contains irreducible submodules isomorphic to the Specht module $\mathcal{S}^{\mu}$ if and only if $\mu = \pi^k$, where π is such that $0 \leqslant |\pi| \leqslant |\mathcal{P}|$, and these are the only irreducible submodules of $\mathcal{V}_{k,\mathcal{P}}$.*

3. Dimensions and multiplicities of the Specht submodules

Given a bijective λ-tableau t on $\{1, 2, \ldots, k\}$ and $\sigma \in \mathrm{Sym}[\lambda]$, we have another bijective λ-tableau $t\sigma$, and the associated $f_{t\sigma} \in \mathcal{M}^{\lambda}$. Define $e_t \in \mathcal{M}^{\lambda}$ as follows:

$$e_t = \sum_{\gamma \in C_\lambda} \mathrm{sign}(\gamma) f_{t\gamma} = \sum_{\gamma \in C_\lambda} \sum_{\rho \in R_\lambda} \mathrm{sign}(\gamma) t\gamma\rho.$$

For example, let $\lambda = (2, 1)$ and $t = \begin{smallmatrix} 1 & 2 \\ 3 & \end{smallmatrix}$. Then $R_\lambda = \{id, \alpha\}$, where α switches the cells in the top row, and $C_\lambda = \{id, \beta\}$, where β switches the cells in the first column. So

$$e_t = f_t - f_{t\beta} = \begin{matrix} 1 & 2 \\ 3 & \end{matrix} + \begin{matrix} 2 & 1 \\ 3 & \end{matrix} - \begin{matrix} 3 & 2 \\ 1 & \end{matrix} - \begin{matrix} 2 & 3 \\ 1 & \end{matrix}.$$

It is easy to check that our definition of e_t is equivalent to the more usual one [11, 2.3.2], where it is called a *polytabloid*:

$$e_t = \kappa_t\{t\}, \quad \text{where} \quad \kappa_t = \sum_{\rho \in C_t} \mathrm{sign}(\rho)\rho \quad \in \mathbb{C}\mathrm{Sym}_k$$

and C_t is the subgroup of Sym_k given by $\{t\gamma t^{-1} \mid \gamma \in C_\lambda\}$.

A λ-tableau t is said to be *standard* if the values assigned by t increase along each row and down each column of $[\lambda]$. In particular, a standard tableau is bijective. The fundamental result on the structure of the Specht modules $\mathcal{S}^{\lambda}$ is as follows [11, 2.5.2].

Lemma 2. *The set of e_t such that t is a standard λ-tableau on $\{1, 2, \ldots, k\}$ is a basis of a submodule of $\mathcal{M}^{\lambda}$ isomorphic to $\mathcal{S}^{\lambda}$.*

It follows from Lemma 2 that the dimension $d(\pi^k)$ of a Specht module $\mathcal{S}^{\pi^k}$ is equal to the number of standard bijections $[\pi^k] \to \{1, 2, \ldots, k\}$. A simple formula for this number can be derived from the well-known *hook formula* [11, 3.10.2]. Given a partition μ and a cell $(i, j) \in [\mu]$, there corresponds a 'hook' consisting of the cells (i, y) with $y \geqslant j$ and the cells (x, j) with $x \geqslant i$. The number of such cells is the *hook-length*

$$h_{ij}(\mu) = (\mu_i - j) + (\mu'_j - i) + 1,$$

where μ'_j is the number of cells in the jth column of μ (that is, the jth part of the conjugate partition μ'). The hook formula for the dimension of $\mathcal{S}^{\mu}$ is

$$d(\mu) = \frac{|\mu|!}{h(\mu)}, \quad \text{where} \quad h(\mu) = \prod_{i,j} h_{ij}(\mu).$$

Lemma 3. *If $|\pi| = \ell$, and π^k is as in Section 2, then*

$$d(\pi^k) = \frac{d(\pi)}{\ell!} \prod_{1 \leqslant i \leqslant \ell} (k - \ell - \pi_i + i).$$

Proof. By the hook formula, it is enough to prove that

$$h(\pi^k) = h(\pi) \left(\frac{k!}{G}\right), \quad \text{where } G = \prod_{1 \leqslant i \leqslant \ell} (k - \ell - \pi_i + i).$$

Since the diagram for π^k is that for π with an extra row, $h(\pi^k) = h(\pi)H$, where H is the product of the hook-lengths corresponding to cells in the top row of π^k. We have to prove that $GH = k!$.

The hook-length corresponding to cell $(0, j)$ is

$$(k - \ell - j + 1) + \pi'_j \quad (1 \leqslant j \leqslant k - \ell),$$

and so H is the product of these numbers. An elementary result [9, p. 3] asserts that, for any partition v and any $m \geqslant v_1$, $n \geqslant v'_1$, the numbers

$$v_j + n + 1 - j \quad (1 \leqslant j \leqslant n) \quad \text{and} \quad n + i - v'_i \quad (1 \leqslant i \leqslant m)$$

are a rearrangement of $1, 2, \ldots, m + n$. Applying this result with $v = \pi'$, $m = \ell$, and $n = k - \ell$ it follows that the numbers

$$(k - \ell - j + 1) + \pi'_j \quad (1 \leqslant j \leqslant k - \ell) \quad \text{and} \quad k - \ell + i - \pi_i \quad (1 \leqslant i \leqslant \ell)$$

are a rearrangement of $1, 2, \ldots, k$. The product of the first set is H and the product of the second set is G, so $GH = k!$ as claimed. □

In terms of the strictly decreasing partition σ of $\frac{1}{2}\ell(\ell + 1)$ associated with π by the rule $\sigma_i = \pi_i + \ell - i$ $(1 \leqslant i \leqslant \ell)$, the preceding result can be written in the form

$$d(\pi^k) = (d(\pi)/\ell!)\ (k - \sigma_1)(k - \sigma_2) \ldots (k - \sigma_\ell).$$

This is clearly a polynomial in k of degree ℓ, and the fact that it takes integer values for all integers k is worth noting.

Lemma 4. *The number of submodules of $\mathcal{V}_k(b)$ isomorphic to $\mathcal{S}^{\pi^k}$ is independent of k and is given by the formula*

$$e(\pi) = \binom{b}{|\pi|} d(\pi).$$

Proof. It follows from Theorem 3 that the required number is equal to the number of semistandard π^k-tableaux on $V \cup \{0\}$, of type $(k - b, 1^b)$. In other words, it is the number of ways of assigning the numbers $0, 1, 2, \ldots, b$ to $[\pi^k]$ in such a way that (i) 0 occurs $k - b$

times and each $i \neq 0$ occurs once, and (ii) the numbers increase weakly in each row and strongly in each column.

In order to satisfy condition (ii), the $k-b$ 0's must be assigned to the first $k-b$ cells of the top row r_0. Let $\ell = |\pi|$, and suppose we have chosen a subset L of size ℓ from $\{1, 2, \ldots, b\}$. Then we can put the elements of L into rows $r_1, r_2, \ldots,$ of $[\pi^k]$, forming a standard π-tableau on L, and the rest (in numerical order) in the last $b-\ell$ cells of r_0. Hence the required number is $\binom{b}{\ell}$ times the number of standard π-tableau on L, and the second term is clearly the same as the number of standard π-tableau on $\{1, 2, \ldots, \ell\}$, that is, $d(\pi)$. $\square$

We shall refer to $e(\pi)$ as the *multiplicity* of $\mathcal{S}^{\pi^k}$.

4. The link with graph colourings

We now focus on the situation when the base graph B is a complete graph K_b with vertex-set $V = \{1, 2, \ldots, b\}$. It follows from the general theory outlined at the beginning of Section 2 that, in principle, the general case can be reduced to this one (for more details, see [10]).

We shall write $\mathcal{V}_k(b)$ for $\mathcal{V}_k(K_b)$. Since there is only one colour-partition of K_b, the trivial one in which each part is a single vertex, $\mathcal{V}_k(b)$ is isomorphic to a single $\mathcal{M}^\lambda$

$$\mathcal{V}_k(b) \approx \mathcal{M}^{\lambda_{k,b}}, \quad \text{where} \quad \lambda_{k,b} = (k-b, 1^b).$$

Our first task is to construct the submodules of $\mathcal{V}_k(b)$ that correspond to the Specht modules. From Lemma 1, we know that these are of the form $\mathcal{S}^{\pi^k}$, where π is any partition such that $0 \leqslant |\pi| \leqslant b$.

Given an injection $F : V \to [\pi^k]$, define $F^* : [\pi^k] \to V \cup \{0\}$ such that F^* is the inverse of F on $\operatorname{Im} F$ and F^* is 0 on all cells not in $\operatorname{Im} F$. In the usual terminology [11, 2.9.1], F^* is a π^k-tableau of *type* $(k-b, 1^b)$. For example, let $k = 10$ and suppose π is the partition $(2, 2, 1)$ of 5. If $b = 6$, we could choose injections $F : \{1, 2, 3, 4, 5, 6\} \to [\pi^k]$ to give the following π^k-tableaux F^*, of type $(4, 1^6)$:

$$\begin{array}{lllll} 0 & 0 & 0 & 0 & 4 \\ 1 & 2 & & & \\ 3 & 6 & & & \\ 5 & & & & \end{array} \qquad \begin{array}{lllll} 0 & 1 & 0 & 0 & 0 \\ 2 & 3 & & & \\ 4 & 6 & & & \\ 5 & & & & \end{array}.$$

Such a tableau is said to be *semistandard* [11, 2.9.5] if the entries increase strictly down each column and weakly along each row of $[\pi^k]$. The first example displayed above is semistandard, but the second is not. Observe that in a semistandard tableau all the $k-b$ zeros occur in the first $k-b$ cells in the top row, and that the restriction of F^* to $[\pi]$ is a standard π-tableau on a subset of V.

The link with k-colourings of K_b can now be made. Given an injection $F : V \to [\pi^k]$, a permutation $\omega \in \operatorname{Sym}[\pi^k]$, and a bijective π^k-tableau t on $\{1, 2, \ldots, k\}$, the composite function $t\omega F$ is such a colouring. So, if we define $f_t F$ and $e_t F$ in the obvious way:

$$f_t F = \sum_{\rho \in R_{\pi^k}} t\rho F, \quad e_t F = \sum_{\gamma \in C_{\pi^k}} \mathrm{sign}(\gamma) f_{t\gamma} F \ ,$$

these are linear combinations of colourings with coefficients ± 1 and thus elements of $\mathcal{V}_k(b)$. Comparison with [11, 2.10.1] gives the fundamental result on the Specht submodules of $\mathcal{V}_k(b)$.

Theorem 5. *For each injection $F : V \to [\pi^k]$, such that F^* is semistandard of type $(k-b, 1^b)$, the set*

$$\{e_t F \mid t \text{ is a standard } \pi^k\text{−tableau on } \{1, 2, \ldots, k\}\}$$

is a basis for a submodule $\mathcal{U}^F$ of $\mathcal{V}_k(b)$, isomorphic to the Specht module $\mathcal{S}^{\pi^k}$. The set of all such $\mathcal{U}^F$ is the complete set of non-identical, irreducible submodules of $\mathcal{V}_k(b)$ that are isomorphic to $\mathcal{S}^{\pi^k}$. □

For a given π, we denote the direct sum of these submodules $\mathcal{U}^F$ of $\mathcal{V}_k(b)$ by $\mathcal{W}^\pi$. That is

$$\mathcal{W}^\pi = \bigoplus \{\mathcal{U}^F \mid F^* \text{ is a semistandard } \pi^k\text{−tableau of type } (k-b, 1^b)\}.$$

Then we have

$$\mathcal{V}_k(b) = \bigoplus \{\mathcal{W}^\pi \mid 0 \leqslant |\pi| \leqslant b\}.$$

5. The chromatic polynomials of bracelets

In this section we shall explain how the decomposition of $\mathcal{V}_k(b)$ into its irreducible submodules leads to explicit formulae for the chromatic polynomials of certain families of graphs. The generalization to $\mathcal{V}_k(B)$ is possible [10] but it will not be discussed here.

We continue to denote the vertex-set of K_b by $V = \{1, 2, \ldots, b\}$. Given a set $L \subseteq V \times V$ and an integer $n \geqslant 3$, we construct the *bracelet* $B_n(b, L)$ as follows. Take n disjoint copies of K_b and link them so that, for each pair $(v, w) \in L$, the vertex v in one copy of K_b is joined to the vertex w in the next copy, with the convention that $n + 1 = 1$. We obtain a ring of n copies of K_b linked by edges in the manner prescribed.

A pair (α, β) of k-colourings of K_b is *compatible with* L if:

$$(v, w) \in L \quad \Longrightarrow \quad \alpha(v) \neq \beta(w).$$

This means that if one copy of K_b is coloured according to α, a second copy of K_b according to β, and they are linked according to L, the resulting graph is properly k-coloured by α and β. The *compatibility matrix* T_L is the matrix whose rows and columns correspond to the k-colourings of K_b, with entries

$$(T_L)_{\alpha\beta} = \begin{cases} 1 & if\,(\alpha, \beta) \text{is compatible with } L; \\ 0 & \text{otherwise.} \end{cases}$$

Note that T_L depends on k, specifically because the number of its rows (and columns) is equal to the number of k-colourings of K_b, the dimension of $\mathcal{V}_k(b)$. Indeed, we can regard T_L as an operator on $\mathcal{V}_k(b)$ in the standard way: if the k-colouring β is identified with an element of $\mathcal{V}_k(b)$, then

$$T_L(\beta) = \sum_{\alpha} (T_L)_{\alpha\beta}\, \alpha = \sum_{\alpha \in L(\beta)} \alpha,$$

where $L(\beta)$ is the set of α such that (α, β) is compatible with L.

The connection between the chromatic polynomial $P(B_n(b, L); k)$ and T_L is given by the following well-known simple result [2].

Lemma 6. *The number of k-colourings of $B_n(b, L)$ is equal to the trace of $(T_L)^n$.*

The symmetric group Sym_k acts on the k-colourings of K_b by permuting the colours. Given $\omega \in \mathrm{Sym}_k$, let

$$(A(\omega))_{\alpha\beta} = \begin{cases} 1 & \text{if} \quad \omega\beta = \alpha, \\ 0 & \text{otherwise.} \end{cases}$$

In other words, A is the matrix representation afforded by the $\mathbb{C}\mathrm{Sym}_k$-module $\mathcal{V}_k(b)$. Recall that the submodule $\mathcal{W}^\pi$ of $\mathcal{V}_k(b)$ is the sum of Specht submodules

$$\mathcal{W}^\pi \; = \; \mathcal{U}^{F_1} \oplus \mathcal{U}^{F_2} \oplus \cdots \oplus \mathcal{U}^{F_{n_\pi}},$$

where $n_\pi = e(\pi) = \binom{b}{|\pi|} d(\pi)$. Let $t_1, t_2, \ldots, t_{m_\pi}$ be the standard π^k-tableau on the set $\{1, 2, \ldots, k\}$, where $m_\pi = m_\pi(k) = d(\pi^k)$. According to Theorem 3, a basis for $\mathcal{U}^{F_j}$ is the set

$$\{e_{t_i} F_j \mid i = 1, 2, \ldots, m_\pi\}.$$

Thus, by changing to the basis $\{e_{t_i} F_j\}$ for each $\mathcal{W}^\pi$, $A(\omega)$ can be reduced to block-diagonal form, with the blocks on the diagonal being matrices of size $m_\pi \times m_\pi$.

Now, it can easily be checked that the action of Sym_k preserves compatibility. In matrix terms, we have

$$T_L\, A(\omega) = A(\omega)\, T_L \quad \text{for all } \omega \in \mathrm{Sym}_k,$$

which means that T_L belongs to the *commutant algebra* of the representation A. For $i = 1, 2, \ldots, m_\pi$, denote the subspace of $\mathcal{W}^\pi$ with basis

$$\{e_{t_i} F_j \mid j = 1, 2, \ldots, n_\pi\}$$

by $\mathcal{Y}^{t_i}$ (Note that this is not a $\mathbb{C}\,\mathrm{Sym}_k$-submodule.) However,

$$\mathcal{W}^\pi \; = \; \mathcal{Y}^{t_1} \oplus \mathcal{Y}^{t_2} \oplus \cdots \oplus \mathcal{Y}^{t_{m_\pi}}$$

and applying Schur's Lemma [11, Sections 1.6 and 1.7] we conclude that, since T_L commutes with $A(\omega)$ for all $\omega \in \mathrm{Sym}_k$, it can be reduced to the form

$$T_L \;\approx \bigoplus_{0 \leqslant |\pi| \leqslant b} I_{m_\pi} \otimes N_L^\pi.$$

Here I_{m_π} is the identity matrix of size m_π and N_L^π is a matrix of size n_π, representing the action of T_L on any one of the subspaces $\mathcal{Y}^{t_i}$. Note that since $n_\pi = e(\pi) = \binom{b}{|\pi|} d(\pi)$, the size of N_L^π does not depend on k, although its entries do.

The explicit formula for $d(\pi^k)$ obtained in Section 4 shows that it can be written as a polynomial in k

$$m_\pi(k) \;=\; d(\pi^k) \;=\; \frac{d(\pi)}{|\pi|!} \prod_{i=1}^{|\pi|} (k - \sigma_i(\pi)),$$

where $\sigma_i(\pi) = \pi_i + |\pi| - i$. Finally, applying the trace formula for the number of colourings (Lemma 6), we have the key result.

Theorem 7. *Suppose integers b and k are given, with $k \geqslant 2b$. For each partition π with $0 \leqslant |\pi| \leqslant b$ let $d(\pi)$ be the dimension of the Specht module $\mathcal{S}^\pi$, and let $m_\pi(k)$ be the polynomial displayed above. Then for any linking set L the number of k-colourings of $B_n(b, L)$ is equal to*

$$\sum_\pi m_\pi(k) \operatorname{tr}\left(N_L^\pi\right)^n,$$

where N_L^π is a matrix of size $\binom{b}{|\pi|} d(\pi)$. □

For example, the number of proper k-colourings of $B_n(3, L)$ for any linking set L can be written as

$$\begin{aligned}
&\operatorname{tr}(N_L^o)^n \;+\; (k-1)\operatorname{tr}(N_L^{(1)})^n \\
&\quad + \frac{1}{2}\,k(k-3)\operatorname{tr}(N_L^{(2)})^n \;+\; \frac{1}{2}\,(k-1)(k-2)\operatorname{tr}(N_L^{(1^2)})^n \\
&\quad + \frac{1}{6}\,k(k-1)(k-5)\operatorname{tr}(N_L^{(3)})^n \\
&\quad + \frac{2}{6}\,k(k-2)(k-4)\operatorname{tr}(N_L^{(2,1)})^n \\
&\quad + \frac{1}{6}\,(k-1)(k-2)(k-3)\operatorname{tr}(N_L^{(1^3)})^n.
\end{aligned}$$

The sizes of the matrices N_L^π are as follows.

π	o	(1)	(2)	(1^2)	(3)	$(2,1)$	(1^3)
size of N_L^π	1×1	3×3	3×3	3×3	1×1	2×2	1×1

Of course, the entries of the matrices N_L^π depend on L, and they are functions of k. It turns out these functions are polynomials, and our next task is to explain how to compute them. The point of the theory developed above is that we can do this by choosing a fixed t and considering the action on the basis elements

$$e_t F_1, e_t F_2, \ldots, e_t F_{n_\pi},$$

where n_π is independent of k.

6. More about the basis elements

Consider a typical basis element $e_t F$. By definition, it is a linear combination of terms of the form $f_u F$, where $u = t\gamma$, $\gamma \in C_{\pi^k}$, and $f_u F$ is a formal sum of colourings

$$\sum_{\rho \in R_{\pi^k}} u\rho F.$$

Lemma 8. *Consider $[\pi]$ as a subset of $[\pi^k]$ in the obvious way, and let $V_F = F^{-1}[\pi]$. Then the colourings that occur in the sum $f_u F$ are just those that agree on V_F with $u\omega F$, for some $\omega \in R_\pi$, and each such colouring occurs $(k-b)!$ times.*

Proof. The row stabilizer R_{π^k} is $\mathrm{Sym}(r_0) \times R_\pi$, so each $\rho \in R_{\pi^k}$ can be written as $\omega\sigma$ with $\omega \in R_\pi$ and $\sigma \in \mathrm{Sym}(r_0)$. Thus we can write

$$f_u F = \sum_{\omega \in R_\pi} \sum_{\sigma \in \mathrm{Sym}(r_0)} u\omega\sigma F.$$

For a fixed ω, each colouring $u\omega\sigma F$ agrees with $u\omega F$ on V_F. Conversely, recall that precisely the last $b - |\pi|$ cells of r_0 belong to $\mathrm{Im}\, F$. Hence if σ fixes these cells pointwise, $\sigma F = F$. The remaining cells of r_0 are $(k - |\pi|) - (b - |\pi|) = k - b$ in number, hence there are $(k-b)!$ colourings $u\omega\sigma F$ that agree with $u\omega F$ on V_F. □

Let X be a subset of the vertex-set V and c an injection from X to $\{1, 2, \ldots, k\}$. We define $\{X \mid c\}$ to be the set of those k-colourings of K_b that agree with c on X. The element of $\mathcal{V}_k(b)$ that is the formal sum of these colourings will be denoted by

$$[X \mid c] = \sum_{c' \in \{X|c\}} c'.$$

In actual calculations (see below) it is often convenient to employ a more explicit form of this notation. If the members of X are listed in order, $x_1, x_2, \ldots$, and $c_1, c_2, \ldots,$ are colours, we write $[x_1, x_2, \ldots \mid c_1, c_2 \ldots]$ for the formal sum of the colourings c' that satisfy $c'(x_1) = c_1,\ c'(x_2) = c_2, \ldots$.

With this notation, the result of Lemma 8 can be written as

$$f_{t\gamma} F = (k-b)! \sum_{\omega \in R_\pi} [V_F \mid t\gamma\omega F]$$

and consequently

$$e_t F = (k-b)! \sum_{\gamma \in C_{\pi^k}} \mathrm{sign}(\gamma) \sum_{\omega \in R_\pi} [V_F \mid t\gamma\omega F].$$

Thus $e_t F$ is expressed as a linear combination of elements of the form $[V_F \mid uF]$. The factor $(k-b)!$ is unimportant, because it is the same for all π.

As an example we calculate explicit basis elements for some typical subspaces $\mathcal{Y}^t$ of $\mathcal{V}_k(b)$, generalizing results formerly obtained by ad hoc methods. The complete calculation for $b = 3$ may be found in [10].

When $\ell = 0$ there is only one partition, the empty partition o, and $o^k = (k)$. There is only one standard o^k-tableau

$$t = \; 1 \; 2 \; \cdots \; k \, .$$

The column stabilizer is trivial, so $e_t = f_t$. There is only one relevant $F : V \to [o^k]$, which corresponds to the semistandard $[o^k]$-tableau of type $(k-b, 1^b)$

$$F^* = \; 0 \; 0 \; \cdots \; 0 \; 1 \; \cdots \; b \, .$$

It follows that $\mathcal{W}^o = \mathcal{Y}^t$ and $\mathcal{Y}^t$ has a basis consisting of one element $e_t F = f_t F$. Here $V_F = \emptyset$, so by Lemma 8, $f_t F = (k-b)!\,[\emptyset \mid tF]$. Since $[\emptyset \mid tF]$ is the formal sum of all colourings, $\mathcal{W}^o$ is the one-dimensional submodule of $\mathcal{V}_k(b)$ spanned by this element.

When $\ell = 1$ there is only one partition, (1), and $(1)^k = (k-1, 1)$. There are $k-1$ standard $(k-1,1)$-tableaux, since the number in the bottom row can be any number r such that $1 < r \leqslant k$:

$$t = \begin{array}{l} 1 \; * \; * \; \cdots \; * \\ r \end{array},$$

where the $*$'s denote the elements of $\{2, \ldots, k\} \setminus \{r\}$ in increasing order. The column stabilizer is $\{id, \beta\}$, where β switches the cells in the first column. Hence

$$e_t = f_t - f_{t\beta}.$$

There are b injections $F_j : \{1, 2, \ldots, b\} \to [(k-1, 1)]$, corresponding to the semistandard $(k-1, 1)$-tableaux of type $(k-b, 1^b)$:

$$F_j^* = \begin{array}{l} 0 \; 0 \; \cdots \; 0 \; * \; * \; \cdots \; * \\ j \end{array},$$

where the $*$'s denote the elements of $V \setminus \{j\}$ in increasing order. We have

$$V_{F_j} = \{j\}, \quad tF_j(j) = r, \quad t\beta F_j(j) = 1.$$

Hence, by Lemma 8,

$$f_t F_j = (k-b)!\,[V_{F_j} \mid tF_j] = (k-b)!\,[j \mid r],$$

$$f_{t\beta} F_j = (k-b)!\,[V_{F_j} \mid t\beta F_j] = (k-b)!\,[j \mid 1]$$

and

$$e_t F_j \;=\; f_t F_j - f_{t\beta} F_j \;=\; (k-b)!\,([j \mid r] - [j \mid 1]).$$

Thus the subspace $\mathcal{Y}^t$ has the basis

$$\{[j \mid r] - [j \mid 1] \mid j = 1, 2, \ldots, b\}.$$

$\mathcal{W}^{(1)}$ is the sum of $(k-1)$ such b-dimensional subspaces, one for each $r \in \{2, 3, \ldots, k\}$.

When $\ell = 2$ there are two partitions, (2) and (1^2). The calculations are similar to those given above, but obviously more complicated. For the partition (2), it turns out that there are $\frac{1}{2}k(k-3)$ standard $(k-2, 2)$-tableaux, one for each pair (r, s) satisfying $1 < r < s \leqslant k$ except (2, 3). Thus $\mathcal{W}^{(2)}$ is the sum of $\frac{1}{2}k(k-3)$ subspaces $\mathcal{Y}^t$. Each has a basis of $\frac{1}{2}b(b-1)$ elements, and when $r > 2$ the basis elements are

$$\begin{aligned}&[i, j \mid r, s] - [i, j \mid 1, s] - [i, j \mid r, 1] + [i, j \mid 1, 2]\\&\quad + [i, j \mid s, r] - [i, j \mid s, 1] - [i, j \mid 1, r] + [i, j \mid 2, 1]\end{aligned}$$

for each unordered pair of vertices $\{i, j\}$ (When $r = 2$ the tableau has a slightly different form, and consequently the basis elements too are different.)

7. The matrices S_M

The key result concerning the matrix T_L is its decomposition in terms of matrices N_L^π (Section 5). In this section, we introduce a set of matrices S_M that will simplify the calculation of N_L^π, for all linking sets L.

We say that $M \subseteq V \times V$ is a *matching* if, given $v, w \in V$, there is at most one pair (v, v') in M, and at most one pair (w', w) in M. The matrix S_M is the matrix whose rows and columns correspond to the k-colourings of K_b, with entries

$$(S_M)_{\alpha\beta} = \begin{cases} 1 & \text{if} \quad (v, w) \in M \ \Rightarrow \ \alpha(v) = \beta(w);\\ 0 & \text{otherwise.}\end{cases}$$

S_M can be regarded as an operator on $\mathcal{V}_k(b)$ in the same way as T_L. In fact, we can describe its action very simply. Given a matching $M \subseteq V \times V$ let M_1, M_2 denote the projections on the factors, and $\mu : M_1 \to M_2$ the bijection such that M is the subset of $V \times V$ consisting of the pairs $(v, \mu(v))$ for all $v \in M_1$. With this notation,

$$S_M(\beta) = \sum_\alpha (S_M)_{\alpha\beta}\,\alpha \ = \sum_{\alpha \in \{M_1 \mid \beta\mu\}} \alpha \ = \ [M_1 \mid \beta\mu].$$

A sieve argument gives the relation between T_L and S_M [5, Theorem 3].

Lemma 9. *For any $L \subseteq V \times V$,*

$$T_L \ = \sum_{M \subseteq L} (-1)^{|M|}\, S_M.$$

It is easily verified that S_M commutes with the action of Sym_k on the colourings. Hence, repeating the argument used for T_L in Section 5, it follows that there exist matrices P_M^π of size $e(\pi)$ such that

$$S_M \ \approx \bigoplus_{0 \leqslant |\pi| \leqslant b} I_{d(\pi^k)} \otimes P_M^\pi.$$

Furthermore, it follows from Lemma 9 that

$$N_L^{\pi} = \sum_{M \subseteq L} (-1)^{|M|} \; P_M^{\pi}.$$

The entries of P_M^{π} are given by the action of S_M on the module $\mathcal{W}^{\pi}$, and according to the theory developed in Section 5, it is enough to calculate the action on one subspace $\mathcal{Y}^t$. In other words, the entries of P_M^{π} are the terms $p(F', F)$ such that

$$S_M(e_t F) = \sum_{F'} p(F', F) \; e_t F'.$$

8. Explicit calculation of the terms

Throughout this section we suppose that we are given k, $V = \{1, 2, \ldots, b\}$, and a partition π such that $|\pi| \leqslant b$. The matching M and the standard tableau $t : [\pi^k] \to \{1, 2, \ldots, k\}$ will also be fixed.

In order to calculate the terms $p(F', F)$ it is convenient to use the bijective representation of semistandard tableaux, introduced in Lemma 5. Let $|\pi| = \ell$, let X be an ℓ-subset of V, and let g be a standard π-tableau on $\{1, 2, \ldots, \ell\}$. If we order the elements of X according to the natural order of V, $x_1 < x_2 < \cdots < x_\ell$, then we have a standard π-tableau g_X on X defined by

$$g_X(r, s) = x_{g(r,s)} \quad (r, s) \in [\pi].$$

The elements of $V \setminus X$ are also ordered in the same way, say $w_1 < w_2 < \cdots < w_{b-\ell}$, and we can define $F(X, g) = F : V \to [\pi^k]$ as follows:

$$F(v) = \begin{cases} g^{-1}(i) & \text{if } v = x_i \in X; \\ (0, k - b + j) & \text{if } v = w_j \notin X. \end{cases}$$

Clearly the associated $F^* : [\pi^k] \to V \cup \{0\}$ is a semistandard π^k-tableau of type $(k-b, 1^b)$. For example, suppose $b = 9$ and $\pi = (3, 1)$. If we take $X = \{2, 4, 7, 8\}$ and

$$g \;=\; \begin{matrix} 1 & 2 & 4 \\ 3 & & \end{matrix}$$

then, provided k is large enough, the semistandard tableau associated with $F = F(X, g)$ is

$$F^* \;=\; \begin{matrix} 0 & 0 & 0 & \cdots & 0 & 1 & 3 & 5 & 6 & 9 \\ 2 & 4 & 8 & & & & & & & \\ 7 & & & & & & & & & \end{matrix}\;.$$

Since $(X, g) \mapsto F$ is a bijection for fixed g, we can take as basis elements of $\mathcal{Y}^t$ the elements

$$b_{X,g} = \frac{1}{(k-b)!} \; e_t F(X, g).$$

When $F = F(X, g)$ we have $V_F = X$ and the restriction of F to V_F is g_X^{-1}, so the results in Section 6 imply that

$$b_{X,g} \;=\; \sum_{\gamma} \mathrm{sign}(\gamma) \sum_{\omega} [X \mid t\gamma\omega g_X^{-1}],$$

where the sums are taken over $\gamma \in C_{\pi^k}$ and $\omega \in R_\pi$.

Since $b_{X,g}$ is a linear combination of terms of the form $[X \mid t\gamma\omega g_X^{-1}]$, we require the effect of S_M on a typical element $[X \mid c]$, which can be computed as follows:

$$S_M[X \mid c] \;=\; S_M\left(\sum_{\beta\in\{X|c\}} \beta\right) \;=\; \sum_{\beta\in\{X|c\}} S_M(\beta) \;=\; \sum_{\beta\in\{X|c\}} \; \sum_{\alpha\in\{M_1|\beta\mu\}} \alpha.$$

By rearranging the double sum and applying another sieve argument, we can obtain [5, Theorem 5] a linear combination of elements of the form $[Y \mid d]$. The explicit form of this result is as follows.

Lemma 10. *A term* $[Y \mid d]$ *occurs in* $S_M[X \mid c]$ *if and only if*
(i) $\mu^{-1}(X \cap M_2) \subseteq Y \subseteq M_1$, *and*
(ii) $d(Y) \subseteq c(X)$, *and whenever* $(y, x) \in M$ *with* $y \in Y$ *and* $xinX$, *then* $d(y) = c(x)$.
If the conditions (i) *and* (ii) *are satisfied the coefficient of* $[Y \mid d]$ *is*

$$(-1)^{|Y|-|X\cap M_2|}\, q(|X \cup M_2|),$$

where $q(s)$ *is the 'falling factorial'* $\langle k - s\rangle_{b-s} = (k - s)(k - s - 1)\ldots(k - b + 1)$.

Note that condition (ii) is equivalent to saying that there is an injection $\theta : Y \to X$ such that $d = c\theta$, and $\theta(y) = \mu(y)$ whenever $\mu(y) \in X$. It follows that $S_M[X \mid t\gamma\omega g_X^{-1}]$ is a linear combination of terms $[Y \mid t\gamma\omega g_X^{-1}\theta]$, where $|Y| \leqslant |X|$. Since S_M leaves invariant each subspace $\mathcal{Y}^t$, when we extend by linearity to $S_M(b_{X,g})$, all terms with $|Y| < |X|$ disappear (a fact which can also be proved directly [10, Theorem 3.10]). This fact is the justification for using the Specht basis elements $b_{X,g}$, rather than the elements $[X \mid c]$, as was done previously [3].

When $\ell = |\pi|$ there is a natural action of Sym_ℓ on the elements e_g, where g is any bijective π-tableau on $\{1, 2, \ldots, \ell\}$, defined by $\sigma * e_g = e_{\sigma g}$. *Young's natural representation* of Sym_ℓ associated with π is obtained by expressing $e_{\sigma g}$ in terms of the standard basis [11, p. 74]:

$$\sigma * e_g = e_{\sigma g} = \sum_h R^\pi_{h,g}(\sigma)\, e_h \quad (g, h \text{ standard}).$$

In the proof of following lemma it will be convenient to define $b_{Y,f}$ by the same explicit formula as that given above for $b_{Y,g}$, whenever the pair (Y, f) is such that f is any bijective (but not necessarily standard) π-tableau on Y.

Lemma 11. *Given* $|\pi|$*-subsets* Y, X *of* V *satisfying condition* (i) *of Lemma* 10, *let* Θ *denote the set of bijections* $Y \to X$ *such that* $\theta(y) = \mu(y)$ *whenever* $\mu(y) \in X$. *For any standard*

π-tableau g on $\{1, 2, \ldots, \ell\}$, the sum

$$\sum_{\theta \in \Theta} \sum_{\gamma} \operatorname{sign}(\gamma) \sum_{\omega} [Y \mid t\gamma\omega g_X^{-1}\theta]$$

is equal to

$$\sum_{\sigma} \sum_{h} R^{\pi}_{h,g}(\sigma^{-1}) b_{Y,h},$$

where the sums are taken over the set of permutations $\sigma \in \mathrm{Sym}_\ell$ such that $\sigma(i) = j$ whenever $(y_i, x_j) \in M$ and the set of standard π-tableaux h.

Proof. We may suppose that Y and X are ordered according to the natural order of V. Then we can associate with a bijection $\theta : Y \to X$ a permutation $\sigma \in \mathrm{Sym}_\ell$, such that

$$\sigma(i) = j \quad \Longleftrightarrow \quad \theta(y_i) = x_j.$$

Under this correspondence $\theta^{-1} g_X$ and $(\sigma^{-1} g)_Y$ define the same π-tableau on Y. Also, taking the sum over bijections $\theta \in \Theta$ is equivalent to taking the sum over the set Σ of permutations $\sigma \in \mathrm{Sym}_\ell$ such that $\sigma(i) = j$ whenever $(y_i, x_j) \in M$. Thus

$$\sum_{\theta \in \Theta} \sum_{\gamma} \operatorname{sign}(\gamma) \sum_{\omega} [Y \mid t\gamma\omega g_X^{-1}\theta] = \sum_{\sigma \in \Sigma} \sum_{\gamma} \operatorname{sign}(\gamma) \sum_{\omega} [Y \mid t\gamma\omega (g^{-1}\sigma)_Y].$$

By definition, the second sum is equal to $\sum_\sigma b_{Y,\sigma^{-1}g}$. Note that $\sigma^{-1}g$ is not generally a standard tableau, and consequently $b_{Y,\sigma^{-1}g}$ is not a basis element. However, it can be expressed as a linear combination of basis elements as follows. Referring to the definitions, the action of Sym_ℓ on the elements $b_{Y,g}$, defined by

$$\tau * b_{Y,g} = b_{Y,\tau g},$$

is the same as the action on the elements e_g. Thus

$$b_{Y,\sigma^{-1}g} = \sigma^{-1} * b_{Y,g} = \sum_{h} R^{\pi}_{h,g}(\sigma^{-1}) b_{Y,h},$$

as claimed. $\square$

Theorem 12. *Suppose the action of S_M on an element $b_{X,g}$ of the basis of $\mathcal{Y}^\ell \subseteq \mathcal{W}^\pi$ is given by*

$$S_M(b_{X,g}) = \sum_{Y,h} p(Y, h; X, g) b_{Y,h}.$$

Then

$$p(Y, h; X, g) = (-1)^{|\pi|} C(Y, X) \sum_{\sigma} R^{\pi}_{h,g}(\sigma^{-1}),$$

where

$C(Y, X) = 0$ *unless* $\mu^{-1}(X \cap M_2) \subseteq Y \subseteq M_1$, *in which case*

$$C(Y, X) = (-1)^{|X \cap M_2|} q(|X \cup M_2|);$$

the sum is taken over all $\sigma \in \mathrm{Sym}_\ell$ *such that* $\sigma(i) = j$ *whenever* $(y_i, x_j) \in M$;
R^π *is Young's natural representation of* Sym_ℓ *associated with* π.

Proof. We have

$$\begin{aligned} S_M(b_{X,g}) &= \sum_\gamma \mathrm{sign}(\gamma) \sum_\omega S_M[X \mid t\gamma\omega g_X^{-1}] \\ &= \sum_\gamma \mathrm{sign}(\gamma) \sum_\omega \sum_{Y,\theta} (-1)^{|Y|-|X\cap M_2|} q(|X \cup M_2|)\, [Y \mid t\gamma\omega g_X^{-1}\theta], \end{aligned}$$

where the last sum is taken over Y and θ such that the conditions of Lemma 10 are satisfied.

Changing the order of summation, and writing $C(Y, X)$ as in the statement of the theorem, we obtain the expression

$$(-1)^{|\pi|} \sum_Y C(Y, X) \sum_{\theta\in\Theta} \sum_\gamma \mathrm{sign}(\gamma) \sum_\omega [Y \mid t\gamma\omega g_X^{-1}\theta].$$

Now it follows from Lemma 11 that

$$\begin{aligned} &\sum_{\theta\in\Theta} \sum_\gamma \mathrm{sign}(\gamma) \sum_\omega [Y \mid t\gamma\omega g_X^{-1}\theta] \\ &\quad = \sum_\sigma \sum_h R^\pi_{h,g}(\sigma^{-1})\, b_{Y,h} \\ &\quad = \sum_h \sum_\sigma R^\pi_{h,g}(\sigma^{-1})\, b_{Y,h}. \qquad \square \end{aligned}$$

The theorem means that we can consider P_M^π as a block matrix with submatrices U_{YX}, where Y, X are $|\pi|$-subsets of V. This submatrix is zero unless Y, X, and M satisfy condition (i) of Lemma 10, in which case U_{YX} has the form

$$\pm\, q(|X \cup M_2|) \sum R^\pi(\sigma^{-1}).$$

This is the 'collapsed' matrix [3], obtained previously by very roundabout arguments.

9. Conclusion

Using the methods described above, the terms involved in the formula for $P(B_n(b, L), k)$, (Theorem 7), can be calculated explicitly and completely for small values of b, and for all L. The polynomials occurring as entries of the matrix P_M^π can be computed once and for all; essentially there is only one calculation for each value of $|M|$ satisfying $|\pi| \leqslant |M| \leqslant b$. Given the catalogue of matrices P_M^π, the matrices N_L^π can be obtained by the sieve formula

(Lemma 9), for any linking set L. The trace of $(N_L^\pi)^n$ is the solution of a linear recursion with coefficients that are polynomials in k (essentially this is Newton's formula applied to the characteristic polynomial).

This approach is followed in [5], where all the matrices P_M^π for $b = 3$ are computed. The Specht bases are not used in that paper, but extensive computations, using the Specht bases, can be found in Reinfeld's thesis [10]; these results are also applicable to the case when the base graph B is not complete. In [5] the matrices N_L^π for the particular linking sets $L = \{11, 22, 33\}$ and $L = \{12, 13, 21, 23, 31, 32\}$ are given, and explicit formulae for the chromatic polynomials of the respective graphs $B_n(3, L)$ obtained. These are 'easy' cases of the formula given in Section 6, in that the eigenvalues of the matrices N_L^π are themselves polynomials in k, and the trace of $(N_L^\pi)^n$ is simply the sum of their nth powers. For example, in the case $L = \{12, 13, 21, 23, 31, 32\}$ the chromatic polynomial is

$$
\begin{aligned}
&(k^3 - 9k^2 + 29k - 32)^n \\
&\quad + (k-1)\,\Big((-2k^2 + 16k - 128)^n + 2(k^2 - 5k + 7)^n\Big) \\
&\quad + \frac{1}{2}\,k(k-3)\,\big((3k-14)^n + 2\big) \\
&\quad + \frac{1}{2}\,(k-1)(k-2)\,\big((k-2)^n + 2(-2k+7)^n\big) \\
&\quad + \frac{1}{6}\,k(k-1)(k-5)\,(-2)^n \\
&\quad + \frac{2}{6}k(k-2)(k-4)\,(2) \\
&\quad + \frac{1}{6}\,(k-1)(k-2)(k-3)\,(-2)^n.
\end{aligned}
$$

It may be worth remarking that although the case $b = 2$ was done by ad hoc methods in 1972, the analogous results for $b = 3$ were not obtained until over 25 years later, and then (initially) by ad hoc methods as well. The situation now is that not only do we have a viable method, but also a theory that explains it, and the prospect of further advances.

What can be said generally about larger values of b, and what happens as $b \to \infty$? In the case when $L = \{11, 22, \ldots, bb\}$, the result for $b = 4$ was given in [3], and some results for larger values of b have been obtained by Chang [6,7]. For certain partitions π, more general results can be obtained. In [3] the terms corresponding to the one-dimensional representations, $\pi = (\ell)$ and $\pi = (1^\ell)$, were obtained explicitly, and for all b. More generally, the arrangement of the terms according to increasing $\ell = |\pi|$ has the property that the terms corresponding to the smallest values of ℓ are in fact the leading terms in the chromatic polynomial. However large b is, the partitions with $0 \leqslant \ell \leqslant r$ determine all the terms of $P(B_n(b, L), k)$ with degree from bn down to $(b-r)n + 1$. Such observations can be used to obtain bounds on the absolute values of the roots of the chromatic polynomials. Last, but not least, the theorem of Beraha et al. [1] can be deployed to analyze the limiting behaviour of the roots as $n \to \infty$ [4].

References

[1] S. Beraha, J. Kahane, N.J. Weiss, Limits of zeros of recursively defined families of polynomials, Stud. Found. Combin. Adv. Math. Suppl. Stud. 1 (1978) 213–232.
[2] N.L. Biggs, Colouring square lattice graphs, Bull. London Math. Soc. 7 (1977) 54–56.
[3] N.L. Biggs, Chromatic polynomials and representations of the symmetric group, Linear Algebra Appl. 356 (2002) 3–26.
[4] N.L. Biggs, Equimodular curves, Discrete Math. 259 (2002) 37–57.
[5] N.L. Biggs, M.H. Klin, P. Reinfeld, Algebraic methods for chromatic polynomials, Europ. J. Combin. 25 (2004) 147–160.
[6] S.C. Chang, Chromatic polynomials for lattice strips with cyclic boundary conditions, Physica A 296 (2001) 495–522.
[7] S.C. Chang, Exact chromatic polynomials of toroidal chains of complete graphs, Physica A 313 (2002) 397.
[8] G.D. James, The Representation Theory of the Symmetric Groups, Lecture Notes in Mathematics, vol. 682, Springer, Berlin, 1978.
[9] I.G. Macdonald, Symmetric Functions and Hall Polynomials, second ed., Clarendon Press, Oxford, 1995.
[10] P. Reinfeld, Algebraic methods for chromatic polynomials, Ph.D. Thesis, University of London, 2003.
[11] B.E. Sagan, The Symmetric Group, Springer, New York, 2001.