

TOPICS IN ALGEBRAIC COMBINATORICS

Richard P. Stanley

VERSION OF 1 FEBRUARY 2013

CONTENTS

Preface		3
Notation		6
Chapter 1	Walks in graphs	9
Chapter 2	Cubes and the Radon transform	21
Chapter 3	Random walks	33
Chapter 4	The Sperner property	45
Chapter 5	Group actions on boolean algebras	59
Chapter 6	Young diagrams and q -binomial coefficients	77
Chapter 7	Enumeration under group action	99
Chapter 8	A glimpse of Young tableaux	131
Appendix	The RSK algorithm	144
Appendix	Plane partitions	147
Chapter 9	The Matrix-Tree Theorem	169
Appendix	Three elegant combinatorial proofs	180
Chapter 10	Eulerian digraphs and oriented trees	189
Chapter 11	Cycles, bonds, and electrical networks	203
11.1	The cycle space and bond space	203
11.2	Bases for the cycle space and bond space	209
11.3	Electrical networks	214
11.4	Planar graphs (sketch)	220
11.5	Squaring the square	223
Chapter 12	Miscellaneous gems of algebraic combinatorics	231
12.1	The 100 prisoners	231
12.2	Oddtown	233

12.3	Complete bipartite partitions of K_n	234
12.4	The nonuniform Fisher inequality	236
12.5	Odd neighborhood covers	238
12.6	Circulant Hadamard matrices	240
12.7	P -recursive functions	246
Hints		257
References		261

Basic Notation

\mathbb{P}	positive integers
\mathbb{N}	nonnegative integers
\mathbb{Z}	integers
\mathbb{Q}	rational numbers
\mathbb{R}	real numbers
\mathbb{C}	complex numbers
$[n]$	the set $\{1, 2, \dots, n\}$ for $n \in \mathbb{N}$ (so $[0] = \emptyset$)
\mathbb{Z}_n	the group of integers modulo n
$R[x]$	the ring of polynomials in the variable x with coefficients in the ring R
Y^X	for sets X and Y , the set of all functions $f: X \rightarrow Y$
$:=$	equal by definition
\mathbb{F}_q	the finite field with q elements
(j)	$1 + q + q^2 + \dots + q^{j-1}$
$\#S$ or $ S $	cardinality (number of elements) of the finite set S
$S \cup T$	the disjoint union of S and T , i.e., $S \cup T$, where $S \cap T = \emptyset$
2^S	the set of all subsets of the set S
$\binom{S}{k}$	the set of k -element subsets of S
$\left(\binom{S}{k}\right)$	the set of k -element multisets on S
KS	the vector space with basis S over the field K

B_n	the poset of all subsets of $[n]$, ordered by inclusion
$\rho(x)$	the rank of the element x in a graded poset
$[x^n]F(x)$	coefficient of x^n in the polynomial or power series $F(x)$
$x \triangleleft y, y \triangleright x$	y covers x in a poset P
δ_{ij}	the Kronecker delta, which equals 1 if $i = j$ and 0 otherwise
$ L $	the sum of the parts (entries) of L , if L is any array of nonnegative integers
$\ell(\lambda)$	length (number of parts) of the partition λ
$p(n)$	number of partitions of the integer $n \geq 0$
$\ker \varphi$	the kernel of a linear transformation or group homomorphism
\mathfrak{S}_n	symmetric group of all permutations of $1, 2, \dots, n$
ι	the identity permutation of a set X , i.e., $\iota(x) = x$ for all $x \in X$

Chapter 1

Walks in graphs


Given a finite set S and integer $k \geq 0$, let $\binom{S}{k}$ denote the set of k -element subsets of S . A *multiset* may be regarded, somewhat informally, as a set with repeated elements, such as $\{1, 1, 3, 4, 4, 4, 6, 6\}$. We are only concerned with how many times each element occurs, and not on any ordering of the elements. Thus for instance $\{2, 1, 2, 4, 1, 2\}$ and $\{1, 1, 2, 2, 2, 4\}$ are the same multiset: they each contain two 1's, three 2's, and one 4 (and no other elements). We say that a multiset M is *on* a set S if every element of M belongs to S . Thus the multiset in the example above is on the set $S = \{1, 3, 4, 6\}$ and also on any set containing S . Let $\left(\binom{S}{k}\right)$ denote the set of k -element multisets on S . For instance, if $S = \{1, 2, 3\}$ then (using abbreviated notation),

$$\binom{S}{2} = \{12, 13, 23\}, \quad \left(\binom{S}{2}\right) = \{11, 22, 33, 12, 13, 23\}.$$

We now define what is meant by a graph. Intuitively, graphs have vertices and edges, where each edge “connects” two vertices (which may be the same). It is possible for two different edges e and e' to connect the same two vertices. We want to be able to distinguish between these two edges, necessitating the following more precise definition. A (finite) *graph* G consists of a *vertex set* $V = \{v_1, \dots, v_p\}$ and *edge set* $E = \{e_1, \dots, e_q\}$, together with a function $\varphi: E \rightarrow \left(\binom{V}{2}\right)$. We think that if $\varphi(e) = uv$ (short for $\{u, v\}$), then e connects u and v or equivalently e is *incident* to u and v . If there is at least one edge incident to u and v then we say that the vertices u and v are *adjacent*. If $\varphi(e) = vv$, then we call e a *loop* at v . If several edges e_1, \dots, e_j ($j > 1$) satisfy $\varphi(e_1) = \dots = \varphi(e_j) = uv$, then we say that there is a *multiple edge*

between u and v . A graph without loops or multiple edges is called *simple*. In this case we can think of E as just a subset of $\binom{V}{2}$ [why?].

The *adjacency matrix* of the graph G is the $p \times p$ matrix $\mathbf{A} = \mathbf{A}(G)$, over the field of complex numbers, whose (i, j) -entry a_{ij} is equal to the number of edges incident to v_i and v_j . Thus \mathbf{A} is a real symmetric matrix (and hence has real eigenvalues) whose trace is the number of loops in G . For instance, if G is the graph



then

$$\mathbf{A}(G) = \begin{bmatrix} 2 & 1 & 0 & 2 & 0 \\ 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 2 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 1 \end{bmatrix}.$$

A *walk* in G of length ℓ from vertex u to vertex v is a sequence $v_1, e_1, v_2, e_2, \dots, v_\ell, e_\ell, v_{\ell+1}$ such that:

- each v_i is a vertex of G
- each e_j is an edge of G
- the vertices of e_i are v_i and v_{i+1} , for $1 \leq i \leq \ell$
- $v_1 = u$ and $v_{\ell+1} = v$.

1.1 Theorem. For any integer $\ell \geq 1$, the (i, j) -entry of the matrix $\mathbf{A}(G)^\ell$ is equal to the number of walks from v_i to v_j in G of length ℓ .

Proof. This is an immediate consequence of the definition of matrix multiplication. Let $\mathbf{A} = (a_{ij})$. The (i, j) -entry of $\mathbf{A}(G)^\ell$ is given by

$$(\mathbf{A}(G)^\ell)_{ij} = \sum a_{ii_1} a_{i_1 i_2} \cdots a_{i_{\ell-1} j},$$

where the sum ranges over all sequences $(i_1, \dots, i_{\ell-1})$ with $1 \leq i_k \leq p$. But since a_{rs} is the number of edges between v_r and v_s , it follows that the summand $a_{ii_1} a_{i_1 i_2} \cdots a_{i_{\ell-1} j}$ in the above sum is just the number (which may be 0) of walks of length ℓ from v_i to v_j of the form

$$v_i, e_1, v_{i_1}, e_2, \dots, v_{i_{\ell-1}}, e_\ell, v_j$$

(since there are a_{ii_1} choices for e_1 , $a_{i_1 i_2}$ choices for e_2 , etc.) Hence summing over all $(i_1, \dots, i_{\ell-1})$ just gives the total number of walks of length ℓ from v_i to v_j , as desired. \square

We wish to use Theorem 1.1 to obtain an explicit formula for the number $(\mathbf{A}(G)^\ell)_{ij}$ of walks of length ℓ in G from v_i to v_j . The formula we give will depend on the eigenvalues of $\mathbf{A}(G)$. The eigenvalues of $\mathbf{A}(G)$ are also called simply the *eigenvalues of G* . Recall that a real symmetric $p \times p$ matrix M has p linearly independent real eigenvectors, which can in fact be chosen to be orthonormal (i.e., orthogonal and of unit length). Let u_1, \dots, u_p be real orthonormal eigenvectors for M , with corresponding eigenvalues $\lambda_1, \dots, \lambda_p$. All vectors u will be regarded as $p \times 1$ *column* vectors, unless specified otherwise. We let t denote transpose, so u^t is a $1 \times p$ *row* vector. Thus the dot (or scalar or inner) product of the vectors u and v is given by $u^t v$ (ordinary matrix multiplication). In particular, $u_i^t u_j = \delta_{ij}$ (the Kronecker delta). Let $U = (u_{ij})$ be the matrix whose columns are u_1, \dots, u_p , denoted $U = [u_1, \dots, u_p]$. Thus U is an orthogonal matrix, so

$$U^t = U^{-1} = \begin{bmatrix} u_1^t \\ \vdots \\ u_p^t \end{bmatrix},$$

the matrix whose rows are u_1^t, \dots, u_p^t . Recall from linear algebra that the matrix U *diagonalizes* M , i.e.,

$$U^{-1} M U = \text{diag}(\lambda_1, \dots, \lambda_p),$$

where $\text{diag}(\lambda_1, \dots, \lambda_p)$ denotes the diagonal matrix with diagonal entries $\lambda_1, \dots, \lambda_p$ (in that order).

1.2 Corollary. *Given the graph G as above, fix the two vertices v_i and v_j . Let $\lambda_1, \dots, \lambda_p$ be the eigenvalues of the adjacency matrix $\mathbf{A}(G)$. Then there exist real numbers c_1, \dots, c_p such that for all $\ell \geq 1$, we have*

$$(\mathbf{A}(G)^\ell)_{ij} = c_1\lambda_1^\ell + \dots + c_p\lambda_p^\ell. \quad (1.1)$$

In fact, if $U = (u_{rs})$ is a real orthogonal matrix such that $U^{-1}\mathbf{A}U = \text{diag}(\lambda_1, \dots, \lambda_p)$, then we have

$$c_k = u_{ik}u_{jk}.$$

Proof. We have [why?]

$$U^{-1}\mathbf{A}^\ell U = \text{diag}(\lambda_1^\ell, \dots, \lambda_p^\ell).$$

Hence

$$\mathbf{A}^\ell = U \cdot \text{diag}(\lambda_1^\ell, \dots, \lambda_p^\ell)U^{-1}.$$

Taking the (i, j) -entry of both sides (and using $U^{-1} = U^t$) gives [why?]

$$(\mathbf{A}^\ell)_{ij} = \sum_k u_{ik}\lambda_k^\ell u_{jk},$$

as desired. □

In order for Corollary 1.2 to be of any use we must be able to compute the eigenvalues $\lambda_1, \dots, \lambda_p$ as well as the diagonalizing matrix U (or eigenvectors u_i). There is one interesting special situation in which it is not necessary to compute U . A *closed walk* in G is a walk that ends where it begins. The number of closed walks in G of length ℓ starting at v_i is therefore given by $(\mathbf{A}(G)^\ell)_{ii}$, so the *total* number $f_G(\ell)$ of closed walks of length ℓ is given by

$$\begin{aligned} f_G(\ell) &= \sum_{i=1}^p (\mathbf{A}(G)^\ell)_{ii} \\ &= \text{tr}(\mathbf{A}(G)^\ell), \end{aligned}$$

where tr denotes trace (sum of the main diagonal entries). Now recall that the trace of a square matrix is the sum of its eigenvalues. If the matrix M has eigenvalues $\lambda_1, \dots, \lambda_p$ then [why?] M^ℓ has eigenvalues $\lambda_1^\ell, \dots, \lambda_p^\ell$. Hence we have proved the following.

1.3 Corollary. *Suppose $\mathbf{A}(G)$ has eigenvalues $\lambda_1, \dots, \lambda_p$. Then the number of closed walks in G of length ℓ is given by*

$$f_G(\ell) = \lambda_1^\ell + \dots + \lambda_p^\ell.$$

We now are in a position to use various tricks and techniques from linear algebra to count walks in graphs. Conversely, it is sometimes possible to count the walks by combinatorial reasoning and use the resulting formula to determine the eigenvalues of G . As a first simple example, we consider the complete graph K_p with vertex set $V = \{v_1, \dots, v_p\}$, and one edge between any two *distinct* vertices. Thus K_p has p vertices and $\binom{p}{2} = \frac{1}{2}p(p-1)$ edges.

1.4 Lemma. *Let J denote the $p \times p$ matrix of all 1's. Then the eigenvalues of J are p (with multiplicity one) and 0 (with multiplicity $p-1$).*

Proof. Since all rows are equal and nonzero, we have $\text{rank}(J) = 1$. Since a $p \times p$ matrix of rank $p-m$ has at least m eigenvalues equal to 0, we conclude that J has at least $p-1$ eigenvalues equal to 0. Since $\text{tr}(J) = p$ and the trace is the sum of the eigenvalues, it follows that the remaining eigenvalue of J is equal to p . \square

1.5 Proposition. *The eigenvalues of the complete graph K_p are as follows: an eigenvalue of -1 with multiplicity $p-1$, and an eigenvalue of $p-1$ with multiplicity one.*

Proof. We have $\mathbf{A}(K_p) = J - I$, where I denotes the $p \times p$ identity matrix. If the eigenvalues of a matrix M are μ_1, \dots, μ_p , then the eigenvalues of $M + cI$ (where c is a scalar) are $\mu_1 + c, \dots, \mu_p + c$ [why?]. The proof follows from Lemma 1.4. \square

1.6 Corollary. *The number of closed walks of length ℓ in K_p from some vertex v_i to itself is given by*

$$(\mathbf{A}(K_p)^\ell)_{ii} = \frac{1}{p}((p-1)^\ell + (p-1)(-1)^\ell). \quad (1.2)$$

(Note that this is also the number of sequences (i_1, \dots, i_ℓ) of numbers $1, 2, \dots, p$ such that $i_1 = i$, no two consecutive terms are equal, and $i_\ell \neq i_1$ [why?].)

Proof. By Corollary 1.3 and Proposition 1.5, the total number of closed walks in K_p of length ℓ is equal to $(p-1)^\ell + (p-1)(-1)^\ell$. By the symmetry of

the graph K_p , the number of closed walks of length ℓ from v_i to itself does not depend on i . (All vertices “look the same.”) Hence we can divide the total number of closed walks by p (the number of vertices) to get the desired answer. \square

A combinatorial proof of Corollary 1.6 is quite tricky (Exercise 1.1). Our algebraic proof gives a first hint of the power of algebra to solve enumerative problems.

What about non-closed walks in K_p ? It’s not hard to diagonalize explicitly the matrix $\mathbf{A}(K_p)$ (or equivalently, to compute its eigenvectors), but there is an even simpler special argument. We have

$$(J - I)^\ell = \sum_{k=0}^{\ell} (-1)^{\ell-k} \binom{\ell}{k} J^k, \quad (1.3)$$

by the binomial theorem.¹ Now for $k > 0$ we have $J^k = p^{k-1}J$ [why?], while $J^0 = I$. (It is not clear *a priori* what is the “correct” value of J^0 , but in order for equation (1.3) to be valid we must take $J^0 = I$.) Hence

$$(J - I)^\ell = \sum_{k=1}^{\ell} (-1)^{\ell-k} \binom{\ell}{k} p^{k-1}J + (-1)^\ell I.$$

Again by the binomial theorem we have

$$(J - I)^\ell = \frac{1}{p}((p-1)^\ell - (-1)^\ell)J + (-1)^\ell I. \quad (1.4)$$

Taking the (i, j) -entry of each side when $i \neq j$ yields

$$(\mathbf{A}(K_p)^\ell)_{ij} = \frac{1}{p}((p-1)^\ell - (-1)^\ell). \quad (1.5)$$

If we take the (i, i) -entry of (1.4) then we recover equation (1.2). Note the curious fact that if $i \neq j$ then

$$(\mathbf{A}(K_p)^\ell)_{ii} - (\mathbf{A}(K_p)^\ell)_{ij} = (-1)^\ell.$$

¹We can apply the binomial theorem in this situation because I and J commute. If A and B are $p \times p$ matrices that don’t necessarily commute, then the best we can say is $(A + B)^2 = A^2 + AB + BA + B^2$, and similarly for higher powers.

We could also have deduced (1.5) from Corollary 1.6 using

$$\sum_{i=1}^p \sum_{j=1}^p (\mathbf{A}(K_p)^\ell)_{ij} = p(p-1)^\ell,$$

the total number of walks of length ℓ in K_p . Details are left to the reader.

We now will show how equation (1.2) itself determines the eigenvalues of $\mathbf{A}(K_p)$. Thus if (1.2) is proved without first computing the eigenvalues of $\mathbf{A}(K_p)$ (which in fact is what we did two paragraphs ago), then we have another means to compute the eigenvalues. The argument we will give can in principle be applied to any graph G , not just K_p . We begin with a simple lemma.

1.7 Lemma. *Suppose $\alpha_1, \dots, \alpha_r$ and β_1, \dots, β_s are nonzero complex numbers such that for all positive integers ℓ , we have*

$$\alpha_1^\ell + \dots + \alpha_r^\ell = \beta_1^\ell + \dots + \beta_s^\ell. \quad (1.6)$$

Then $r = s$ and the α 's are just a permutation of the β 's.

Proof. We will use the powerful method of *generating functions*. Let x be a complex number whose absolute value (or modulus) is close to 0. Multiply (1.6) by x^ℓ and sum on all $\ell \geq 1$. The geometric series we obtain will converge, and we get

$$\frac{\alpha_1 x}{1 - \alpha_1 x} + \dots + \frac{\alpha_r x}{1 - \alpha_r x} = \frac{\beta_1 x}{1 - \beta_1 x} + \dots + \frac{\beta_s x}{1 - \beta_s x}. \quad (1.7)$$

This is an identity valid for sufficiently small (in modulus) complex numbers. By clearing denominators we obtain a polynomial identity. But if two polynomials in x agree for infinitely many values, then they are the same polynomial [why?]. Hence equation (1.7) is actually valid for *all* complex numbers x (ignoring values of x which give rise to a zero denominator).

Fix a complex number $\gamma \neq 0$. Multiply (1.7) by $1 - \gamma x$ and let $x \rightarrow 1/\gamma$. The left-hand side becomes the number of α_i 's which are equal to γ , while the right-hand side becomes the number of β_j 's which are equal to γ [why?]. Hence these numbers agree for all γ , so the lemma is proved. \square

1.8 Example. Suppose that G is a graph with 12 vertices and that the number of closed walks of length ℓ in G is equal to $3 \cdot 5^\ell + 4^\ell + 2(-2)^\ell + 4$. Then it follows from Corollary 1.3 and Lemma 1.7 [why?] that the eigenvalues of $\mathbf{A}(G)$ are given by 5, 5, 5, 4, -2, -2, 1, 1, 1, 1, 0, 0.

Notes for Chapter 1

The connection between graph eigenvalues and the enumeration of walks is considered “folklore.” The subject of *spectral graph theory*, which is concerned with the spectrum (multiset of eigenvalues) of various matrices associated with graphs, began around 1931 in the area of quantum chemistry. The first mathematical paper was published by L. Collatz and U. Sinogowitz in 1957. A good general reference is the book² [26] by Cvetković, Doob, and Sachs. Two textbooks on this subject are by Cvetković, Rowlinson, and Simić [27] and by Brouwer and Haemers [14].

²All citations to the literature refer to the bibliography beginning on page 261.

Chapter 2

Cubes and the Radon transform

Let us now consider a more interesting example of a graph G , one whose eigenvalues have come up in a variety of applications. Let \mathbb{Z}_2 denote the cyclic group of order 2, with elements 0 and 1, and group operation being addition modulo 2. Thus $0 + 0 = 0$, $0 + 1 = 1 + 0 = 1$, $1 + 1 = 0$. Let \mathbb{Z}_2^n denote the direct product of \mathbb{Z}_2 with itself n times, so the elements of \mathbb{Z}_2^n are n -tuples (a_1, \dots, a_n) of 0's and 1's, under the operation of component-wise addition. Define a graph C_n , called the n -cube, as follows: the vertex set of C_n is given by $V(C_n) = \mathbb{Z}_2^n$, and two vertices u and v are connected by an edge if they differ in exactly one component. Equivalently, $u + v$ has exactly one nonzero component. If we regard \mathbb{Z}_2^n as consisting of *real* vectors, then these vectors form the set of vertices of an n -dimensional cube. Moreover, two vertices of the cube lie on an edge (in the usual geometric sense) if and only if they form an edge of C_n . This explains why C_n is called the n -cube. We also see that walks in C_n have a nice geometric interpretation — they are simply walks along the edges of an n -dimensional cube.

We want to determine explicitly the eigenvalues and eigenvectors of C_n . We will do this by a somewhat indirect but extremely useful and powerful technique, the finite Radon transform. Let \mathcal{V} denote the set of all functions $f: \mathbb{Z}_2^n \rightarrow \mathbb{R}$, where \mathbb{R} denotes the field of real numbers.¹ Note that \mathcal{V} is a vector space over \mathbb{R} of dimension 2^n [why?]. If $u = (u_1, \dots, u_n)$ and $v =$

¹For abelian groups other than \mathbb{Z}_2^n it is necessary to use complex numbers rather than real numbers. We could use complex numbers here, but there is no need to do so.

(v_1, \dots, v_n) are elements of \mathbb{Z}_2^n , then define their *dot product* by

$$u \cdot v = u_1v_1 + \dots + u_nv_n, \quad (2.1)$$

where the computation is performed modulo 2. Thus we regard $u \cdot v$ as an element of \mathbb{Z}_2 . The expression $(-1)^{u \cdot v}$ is defined to be the *real number* $+1$ or -1 , depending on whether $u \cdot v = 0$ or 1 , respectively. Since for integers k the value of $(-1)^k$ depends only on $k \pmod{2}$, it follows that we can treat u and v as integer vectors without affecting the value of $(-1)^{u \cdot v}$. Thus, for instance, formulas such as

$$(-1)^{u \cdot (v+w)} = (-1)^{u \cdot v + u \cdot w} = (-1)^{u \cdot v} (-1)^{u \cdot w}$$

are well-defined and valid. From a more algebraic viewpoint, the map $\mathbb{Z} \rightarrow \{-1, 1\}$ sending n to $(-1)^n$ is a group homomorphism, where of course the product on $\{-1, 1\}$ is multiplication.

We now define two important bases of the vector space \mathcal{V} . There will be one basis element of each basis for each $u \in \mathbb{Z}_2^n$. The first basis, denoted B_1 , has elements f_u defined as follows:

$$f_u(v) = \delta_{uv}, \quad (2.2)$$

the Kronecker delta. It is easy to see that B_1 is a basis, since any $g \in \mathcal{V}$ satisfies

$$g = \sum_{u \in \mathbb{Z}_2^n} g(u) f_u \quad (2.3)$$

[why?]. Hence B_1 spans \mathcal{V} , so since $\#B_1 = \dim \mathcal{V} = 2^n$, it follows that B_1 is a basis. The second basis, denoted B_2 , has elements χ_u defined as follows:

$$\chi_u(v) = (-1)^{u \cdot v}.$$

In order to show that B_2 is a basis, we will use an inner product on \mathcal{V} (denoted $\langle \cdot, \cdot \rangle$) defined by

$$\langle f, g \rangle = \sum_{u \in \mathbb{Z}_2^n} f(u)g(u).$$

Note that this inner product is just the usual dot product with respect to the basis B_1 .

2.1 Lemma. *The set $B_2 = \{\chi_u : u \in \mathbb{Z}_2^n\}$ forms a basis for \mathcal{V} .*

Proof. Since $\#B_2 = \dim \mathcal{V} (= 2^n)$, it suffices to show that B_2 is linearly independent. In fact, we will show that the elements of B_2 are orthogonal². We have

$$\begin{aligned} \langle \chi_u, \chi_v \rangle &= \sum_{w \in \mathbb{Z}_2^n} \chi_u(w) \chi_v(w) \\ &= \sum_{w \in \mathbb{Z}_2^n} (-1)^{(u+v) \cdot w}. \end{aligned}$$

It is left as an easy exercise to the reader to show that for any $y \in \mathbb{Z}_2^n$, we have

$$\sum_{w \in \mathbb{Z}_2^n} (-1)^{y \cdot w} = \begin{cases} 2^n, & \text{if } y = \mathbf{0} \\ 0, & \text{otherwise,} \end{cases}$$

where $\mathbf{0}$ denotes the identity element of \mathbb{Z}_2^n (the vector $(0, 0, \dots, 0)$). Thus $\langle \chi_u, \chi_v \rangle = 0$ if and only if $u + v = \mathbf{0}$, i.e., $u = v$, so the elements of B_2 are orthogonal (and nonzero). Hence they are linearly independent as desired. \square

We now come to the key definition of the Radon transform.

Given a subset Γ of \mathbb{Z}_2^n and a function $f \in \mathcal{V}$, define a new function $\Phi_\Gamma f \in \mathcal{V}$ by

$$\Phi_\Gamma f(v) = \sum_{w \in \Gamma} f(v + w).$$

The function $\Phi_\Gamma f$ is called the (*discrete* or *finite*) *Radon transform* of f (on the group \mathbb{Z}_2^n , with respect to the subset Γ).

We have defined a map $\Phi_\Gamma: \mathcal{V} \rightarrow \mathcal{V}$. It is easy to see that Φ_Γ is a linear transformation; we want to compute its eigenvalues and eigenvectors.

2.2 Theorem. *The eigenvectors of Φ_Γ are the functions χ_u , where $u \in \mathbb{Z}_2^n$. The eigenvalue λ_u corresponding to χ_u (i.e., $\Phi_\Gamma \chi_u = \lambda_u \chi_u$) is given by*

$$\lambda_u = \sum_{w \in \Gamma} (-1)^{u \cdot w}.$$

²Recall from linear algebra that nonzero orthogonal vectors in a real vector space are linearly independent.

Proof. Let $v \in \mathbb{Z}_2^n$. Then

$$\begin{aligned}\Phi_\Gamma \chi_u(v) &= \sum_{w \in \Gamma} \chi_u(v+w) \\ &= \sum_{w \in \Gamma} (-1)^{u \cdot (v+w)} \\ &= \left(\sum_{w \in \Gamma} (-1)^{u \cdot w} \right) (-1)^{u \cdot v} \\ &= \left(\sum_{w \in \Gamma} (-1)^{u \cdot w} \right) \chi_u(v).\end{aligned}$$

Hence

$$\Phi_\Gamma \chi_u = \left(\sum_{w \in \Gamma} (-1)^{u \cdot w} \right) \chi_u,$$

as desired. \square

Note that because the χ_u 's form a basis for \mathcal{V} by Lemma 2.1, it follows that Theorem 2.2 yields a complete set of eigenvalues and eigenvectors for Φ_Γ . Note also that the eigenvectors χ_u of Φ_Γ are independent of Γ ; only the eigenvalues depend on Γ .

Now we come to the payoff. Let $\Delta = \{\delta_1, \dots, \delta_n\}$, where δ_i is the i th unit coordinate vector (i.e., δ_i has a 1 in position i and 0's elsewhere). Note that the j th coordinate of δ_i is just δ_{ij} (the Kronecker delta), explaining our notation δ_i . Let $[\Phi_\Delta]$ denote the matrix of the linear transformation $\Phi_\Delta: \mathcal{V} \rightarrow \mathcal{V}$ with respect to the basis B_1 of \mathcal{V} given by (2.2).

2.3 Lemma. *We have $[\Phi_\Delta] = \mathbf{A}(C_n)$, the adjacency matrix of the n -cube.*

Proof. Let $v \in \mathbb{Z}_2^n$. We have

$$\begin{aligned}\Phi_\Delta f_u(v) &= \sum_{w \in \Delta} f_u(v+w) \\ &= \sum_{w \in \Delta} f_{u+w}(v),\end{aligned}$$

since $u = v + w$ if and only if $u + w = v$. There follows

$$\Phi_\Delta f_u = \sum_{w \in \Delta} f_{u+w}. \tag{2.4}$$

Equation (2.4) says that the (u, v) -entry of the matrix Φ_Δ is given by

$$(\Phi_\Delta)_{uv} = \begin{cases} 1, & \text{if } u + v \in \Delta \\ 0, & \text{otherwise.} \end{cases}$$

Now $u + v \in \Delta$ if and only if u and v differ in exactly one coordinate. This is just the condition for uv to be an edge of C_n , so the proof follows. \square

2.4 Corollary. *The eigenvectors E_u ($u \in \mathbb{Z}_2^n$) of $\mathbf{A}(C_n)$ (regarded as linear combinations of the vertices of C_n , i.e., of the elements of \mathbb{Z}_2^n) are given by*

$$E_u = \sum_{v \in \mathbb{Z}_2^n} (-1)^{u \cdot v} v. \quad (2.5)$$

The eigenvalue λ_u corresponding to the eigenvector E_u is given by

$$\lambda_u = n - 2\omega(u), \quad (2.6)$$

where $\omega(u)$ is the number of 1's in u . (The integer $\omega(u)$ is called the Hamming weight or simply the weight of u .) Hence $\mathbf{A}(C_n)$ has $\binom{n}{i}$ eigenvalues equal to $n - 2i$, for each $0 \leq i \leq n$.

Proof. For any function $g \in \mathcal{V}$ we have by (2.3) that

$$g = \sum_v g(v) f_v.$$

Applying this equation to $g = \chi_u$ gives

$$\chi_u = \sum_v \chi_u(v) f_v = \sum_v (-1)^{u \cdot v} f_v. \quad (2.7)$$

Equation (2.7) expresses the eigenvector χ_u of Φ_Δ (or even Φ_Γ for any $\Gamma \subseteq \mathbb{Z}_2^n$) as a linear combination of the functions f_v . But Φ_Δ has the same matrix with respect to the basis of the f_v 's as $\mathbf{A}(C_n)$ has with respect to the vertices v of C_n . Hence the expansion of the eigenvectors of Φ_Δ in terms of the f_v 's has the same coefficients as the expansion of the eigenvectors of $\mathbf{A}(C_n)$ in terms of the v 's, so equation (2.5) follows.

According to Theorem 2.2 the eigenvalue λ_u corresponding to the eigenvector χ_u of Φ_Δ (or equivalently, the eigenvector E_u of $\mathbf{A}(C_n)$) is given by

$$\lambda_u = \sum_{w \in \Delta} (-1)^{u \cdot w}. \quad (2.8)$$

Now $\Delta = \{\delta_1, \dots, \delta_n\}$, and $\delta_i \cdot u$ is 1 if u has a one in its i th coordinate and is 0 otherwise. Hence the sum in (2.8) has $n - \omega(u)$ terms equal to +1 and $\omega(u)$ terms equal to -1, so $\lambda_u = (n - \omega(u)) - \omega(u) = n - 2\omega(u)$, as claimed. \square

We have all the information needed to count walks in C_n .

2.5 Corollary. *Let $u, v \in \mathbb{Z}_2^n$, and suppose that $\omega(u + v) = k$ (i.e., u and v disagree in exactly k coordinates). Then the number of walks of length ℓ in C_n between u and v is given by*

$$(\mathbf{A}^\ell)_{uv} = \frac{1}{2^n} \sum_{i=0}^n \sum_{j=0}^k (-1)^j \binom{k}{j} \binom{n-k}{i-j} (n-2i)^\ell, \quad (2.9)$$

where we set $\binom{n-k}{i-j} = 0$ if $j > i$. In particular,

$$(\mathbf{A}^\ell)_{uu} = \frac{1}{2^n} \sum_{i=0}^n \binom{n}{i} (n-2i)^\ell. \quad (2.10)$$

Proof. Let E_u and λ_u be as in Corollary 2.4. In order to apply Corollary 1.2, we need the eigenvectors to be of *unit* length (where we regard the f_v 's as an orthonormal basis of \mathcal{V}). By equation (2.5), we have

$$|E_u|^2 = \sum_{v \in \mathbb{Z}_2^n} ((-1)^{u \cdot v})^2 = 2^n.$$

Hence we should replace E_u by $E'_u = \frac{1}{2^{n/2}} E_u$ to get an orthonormal basis. According to Corollary 1.2, we thus have

$$(\mathbf{A}^\ell)_{uv} = \frac{1}{2^n} \sum_{w \in \mathbb{Z}_2^n} E_{uw} E_{vw} \lambda_w^\ell.$$

Now E_{uw} by definition is the coefficient of f_w in the expansion (2.5), i.e., $E_{uw} = (-1)^{u \cdot w}$ (and similarly for E_v), while $\lambda_w = n - 2\omega(w)$. Hence

$$(\mathbf{A}^\ell)_{uv} = \frac{1}{2^n} \sum_{w \in \mathbb{Z}_2^n} (-1)^{(u+v) \cdot w} (n - 2\omega(w))^\ell. \quad (2.11)$$

The number of vectors w of Hamming weight i which have j 1's in common with $u + v$ is $\binom{k}{j} \binom{n-k}{i-j}$, since we can choose the j 1's in $u + v$ which agree

with w in $\binom{k}{j}$ ways, while the remaining $i - j$ 1's of w can be inserted in the $n - k$ remaining positions in $\binom{n-k}{i-j}$ ways. Since $(u + v) \cdot w \equiv j \pmod{2}$, the sum (2.11) reduces to (2.9) as desired. Clearly setting $u = v$ in (2.9) yields (2.10), completing the proof. \square

It is possible to give a direct proof of equation (2.10) avoiding linear algebra, though we do not do so here. Thus by Corollary 1.3 and Lemma 1.7 (exactly as was done for K_n) we have another determination of the eigenvalues of C_n . With a little more work one can also obtain a direct proof of equation (2.9). Later in Example 9.12, however, we will use the eigenvalues of C_n to obtain a combinatorial result for which a nonalgebraic proof was found only recently and is by no means easy.

2.6 Example. Setting $k = 1$ in (2.9) yields

$$\begin{aligned} (\mathbf{A}^\ell)_{uv} &= \frac{1}{2^n} \sum_{i=0}^n \left[\binom{n-1}{i} - \binom{n-1}{i-1} \right] (n-2i)^\ell \\ &= \frac{1}{2^n} \sum_{i=0}^{n-1} \binom{n-1}{i} \frac{(n-2i)^{\ell+1}}{n-i}. \end{aligned}$$

NOTE (for those familiar with the representation theory of finite groups). The functions $\chi_u: \mathbb{Z}_2^n \rightarrow \mathbb{R}$ are just the irreducible (complex) characters of the group \mathbb{Z}_2^n , and the orthogonality of the χ_u 's shown in the proof of Lemma 2.1 is the usual orthogonality relation for the irreducible characters of a finite group. The results of this chapter extend readily to any finite abelian group. Exercise 2.5 does the case \mathbb{Z}_n , the cyclic group of order n . For nonabelian finite groups the situation is much more complicated because not all irreducible representations have degree one (i.e., are homomorphisms $G \rightarrow \mathbb{C}^*$, the multiplicative group of \mathbb{C}), and there do not exist formulas as explicit as the ones for abelian groups.

We can give a little taste of the situation for arbitrary groups as follows. Let G be a finite group, and let $\mathbf{M}(G)$ be its multiplication table. Regard the entries of $\mathbf{M}(G)$ as *commuting* indeterminates, so that $\mathbf{M}(G)$ is simply a matrix with indeterminate entries. For instance, let $G = \mathbb{Z}_3$. Let the elements of G be a, b, c , where say a is the identity. Then

$$\mathbf{M}(G) = \begin{bmatrix} a & b & c \\ b & c & a \\ c & a & b \end{bmatrix}.$$

We can compute that $\det \mathbf{M}(G) = (a+b+c)(a+\omega b+\omega^2 c)(a+\omega^2 b+\omega c)$, where $\omega = e^{2\pi i/3}$. In general, when G is abelian, Dedekind knew that $\det \mathbf{M}(G)$ factors into certain explicit linear factors over \mathbb{C} . Theorem 2.2 is equivalent to this statement for the group $G = \mathbb{Z}_2^n$ [why?]. Equation (12.5) gives the factorization for $G = \mathbb{Z}_n$. (For each $w \in G$ one needs to interchange the row indexed by the group element w with the row indexed by w^{-1} in order to convert $\mathbf{M}(\mathbb{Z}_n)$ to the circulant matrices of equation (12.5), but these operations only affect the sign of the determinant.) Dedekind asked Frobenius about the factorization of $\det \mathbf{M}(G)$, known as the *group determinant*, for nonabelian finite G . For instance, let $G = \mathfrak{S}_3$ (the symmetric group of all permutations of 1, 2, 3), with elements (in cycle notation) $a = (1)(2)(3)$, $b = (1, 2)(3)$, $c = (1, 3)(2)$, $d = (1)(2, 3)$, $e = (1, 2, 3)$, $f = (1, 3, 2)$. Then $\det \mathbf{M}(G) = f_1 f_2 f_3^2$, where

$$\begin{aligned} f_1 &= a + b + c + d + e + f \\ f_2 &= -a + b + c + d - e - f \\ f_3 &= a^2 - b^2 - c^2 - d^2 + e^2 + f^2 - ae - af + bc + bd + cd - ef. \end{aligned}$$

Frobenius showed that in general there is a set \mathcal{P} of irreducible homogeneous polynomials f , of some degree d_f , where $\#\mathcal{P}$ is the number of conjugacy classes of G , for which

$$\det \mathbf{M}(G) = \prod_{f \in \mathcal{P}} f^{d_f}.$$

Note that taking the degree of both sides gives $\#G = \sum_f d_f^2$. Frobenius' result was a highlight in his development of group representation theory. The numbers d_f are just the degrees of the irreducible (complex) representations of G . For the symmetric group \mathfrak{S}_n , these degrees are the numbers f^λ of Theorem 8.1, and Appendix 1 to Chapter 8 gives a bijective proof that $\sum_\lambda (f^\lambda)^2 = n!$.

Notes for Chapter 2

The Radon transform first arose in a continuous setting in the paper [92] of J. K. A. Radon and has been applied to such areas as computerized tomography. The finite version was first defined by E. Bolker [10]. For some further applications to combinatorics see J. Kung [68]. For the Radon transform on the n -cube \mathbb{Z}_2^n , see P. Diaconis and R. Graham [29]. For the generalization to \mathbb{Z}_k^n , see M. DeDeo and E. Velasquez [28].

For an exposition of the development of group representation theory by Frobenius and other pioneers, see the survey articles of Hawkins [55][56][57].

Chapter 3

Random walks

Let G be a finite graph. We consider a random walk on the vertices of G of the following type. Start at a vertex u . (The vertex u could be chosen randomly according to some probability distribution or could be specified in advance.) Among all the edges incident to u , choose one uniformly at random (i.e., if there are k edges incident to u , then each of these edges is chosen with probability $1/k$). Travel to the vertex v at the other end of the chosen edge and continue as before from v . Readers with some familiarity with probability theory will recognize this random walk as a special case of a finite-state Markov chain. Many interesting questions may be asked about such walks; the basic one is to determine the probability of being at a given vertex after a given number ℓ of steps.

Suppose vertex u has *degree* d_u , i.e., there are d_u edges incident to u (counting loops at u once only). Let $\mathbf{M} = \mathbf{M}(G)$ be the matrix whose rows and columns are indexed by the vertex set $\{v_1, \dots, v_p\}$ of G , and whose (u, v) -entry is given by

$$\mathbf{M}_{uv} = \frac{\mu_{uv}}{d_u}, \quad (3.1)$$

where μ_{uv} is the number of edges between u and v (which for simple graphs will be 0 or 1). Thus \mathbf{M}_{uv} is just the probability that if one starts at u , then the next step will be to v . We call \mathbf{M} the *probability matrix* associated with G . An elementary probability theory argument (equivalent to Theorem 1.1) shows that if ℓ is a positive integer, then $(\mathbf{M}^\ell)_{uv}$ is equal to the probability that one ends up at vertex v in ℓ steps given that one has started at u . Suppose now that the starting vertex is not specified, but rather we are given probabilities ρ_u summing to 1 and that we start at vertex u with probability

ρ_u . Let P be the row vector $P = [\rho_{v_1}, \dots, \rho_{v_p}]$. Then again an elementary argument shows that if $PM^\ell = [\sigma_{v_1}, \dots, \sigma_{v_p}]$, then σ_v is the probability of ending up at v in ℓ steps (with the given starting distribution). By reasoning as in Chapter 1, we see that if we know the eigenvalues and eigenvectors of \mathbf{M} , then we can compute the crucial probabilities $(\mathbf{M}^\ell)_{uv}$ and σ_u .

Since the matrix \mathbf{M} is not the same as the adjacency matrix \mathbf{A} , what does all this have to do with adjacency matrices? The answer is that in one important case \mathbf{M} is just a scalar multiple of \mathbf{A} . We say that the graph G is *regular of degree d* if each $d_u = d$, i.e., each vertex is incident to d edges. In this case it's easy to see that $\mathbf{M}(G) = \frac{1}{d}\mathbf{A}(G)$. Hence the eigenvectors E_u of $\mathbf{M}(G)$ and $\mathbf{A}(G)$ are the same, and the eigenvalues are related by $\lambda_u(\mathbf{M}) = \frac{1}{d}\lambda_u(\mathbf{A})$. Thus random walks on a regular graph are closely related to the adjacency matrix of the graph.

3.1 Example. Consider a random walk on the n -cube C_n which begins at the “origin” (the vector $(0, \dots, 0)$). What is the probability p_ℓ that after ℓ steps one is again at the origin? Before applying any formulas, note that after an even (respectively, odd) number of steps, one must be at a vertex with an even (respectively, odd) number of 1's. Hence $p_\ell = 0$ if ℓ is odd. Now note that C_n is regular of degree n . Thus by (2.6), we have

$$\lambda_u(\mathbf{M}(C_n)) = \frac{1}{n}(n - 2\omega(u)).$$

By (2.10) we conclude that

$$p_\ell = \frac{1}{2^n n^\ell} \sum_{i=0}^n \binom{n}{i} (n - 2i)^\ell.$$

Note that the above expression for p_ℓ does indeed reduce to 0 when ℓ is odd.

It is worth noting that even though the probability matrix \mathbf{M} need not be a symmetric matrix, nonetheless it has only real eigenvalues.

3.2 Theorem. *Let G be a finite graph. Then the probability matrix $\mathbf{M} = \mathbf{M}(G)$ is diagonalizable and has only real eigenvalues.*

Proof. Since we are assuming that G is connected and has at least two vertices, it follows that $d_v > 0$ for every vertex v of G . Let \mathbf{D} be the diagonal matrix whose rows and columns are indexed by the vertices of G , with

$D_{vv} = \sqrt{d_v}$. Then

$$\begin{aligned} (DMD^{-1})_{uv} &= \sqrt{d_u} \cdot \frac{\mu_{uv}}{d_u} \cdot \frac{1}{\sqrt{d_v}} \\ &= \frac{\mu_{uv}}{\sqrt{d_u d_v}}. \end{aligned}$$

Hence DMD^{-1} is a symmetric matrix and thus has only real eigenvalues. But if B and C are any $p \times p$ matrices with C invertible, then B and CBC^{-1} have the same characteristic polynomial and hence the same eigenvalues. Therefore all the eigenvalues of M are real. Moreover, B is diagonalizable if and only if CBC^{-1} is diagonalizable. (In fact, B and CBC^{-1} have the same Jordan canonical form.) Since a symmetric matrix is diagonalizable, it follows that M is also diagonalizable. \square

Let us give one further example of the connection between linear algebra and random walks on graphs. Let u and v be vertices of a connected graph G . Define the *access time* or *hitting time* $H(u, v)$ to be the expected number of steps that a random walk (as defined above) starting at u takes to reach v for the first time. Thus if the probability is p_n that we reach v for the first time in n steps, then by definition of expectation we have

$$H(u, v) = \sum_{n \geq 1} np_n. \quad (3.2)$$

Conceivably this sum could be infinite, though we will see below that this is not the case. Note that $H(v, v) = 0$.

As an example, suppose that G has three vertices u, v, w with an edge between u and w and another edge between w and v . We can compute $H(u, v)$ as follows. After one step we will be at w . Then with probability $\frac{1}{2}$ we will step to v and with probability $\frac{1}{2}$ back to u . Hence [why?]

$$H(u, v) = \frac{1}{2} \cdot 2 + \frac{1}{2}(2 + H(u, v)). \quad (3.3)$$

Solving this linear equation gives $H(u, v) = 4$.

We want to give a formula for the access time $H(u, v)$ in terms of linear algebra. The proof requires some basic results on eigenvalues and eigenvectors of nonnegative matrices, which we will explain and then state without proof. An $r \times r$ real matrix B is called *nonnegative* if every entry is nonnegative. We say that B is *irreducible* if it is not the 1×1 matrix $[0]$ and if

there does not exist a permutation matrix \mathbf{P} (a matrix with one 1 in every row and column, and all other entries 0) such that

$$\mathbf{PBP}^{-1} = \begin{bmatrix} \mathbf{C} & \mathbf{D} \\ \mathbf{0} & \mathbf{E} \end{bmatrix},$$

where \mathbf{C} and \mathbf{E} are square matrices of size greater than zero. For instance, the adjacency matrix \mathbf{A} and probability matrix \mathbf{M} of a graph G are irreducible if and only if G is connected and is not an *isolated vertex* (that is, a vertex v incident to no edges, not even a loop from v to itself). We now state without proof a version of the *Perron–Frobenius theorem*. There are some other parts of the Perron-Frobenius theorem that we don't need here and are omitted.

3.3 Theorem. *Let \mathbf{B} be a nonnegative irreducible square matrix. If ρ is the maximum absolute value of the eigenvalues of \mathbf{B} , then $\rho > 0$, and there is an eigenvalue equal to ρ . Moreover, there is an eigenvector for ρ (unique up to multiplication by a positive real number) all of whose entries are positive.*

Now let \mathbf{M} be the probability matrix defined by equation (3.1). Let $\mathbf{M}[v]$ denote \mathbf{M} with the row and column indexed by v deleted. Thus if G has p vertices, then $\mathbf{M}[v]$ is a $(p-1) \times (p-1)$ matrix. Let $T[v]$ be the column vector of length $p-1$ whose rows are indexed by the vertices $w \neq v$, with $T[v]_w = \mu(w, v)/d_w$. Write I_{p-1} for the identity matrix of size $p-1$.

3.4 Theorem. *The matrix $I_{p-1} - \mathbf{M}[v]$ is invertible, and*

$$H(u, v) = ((I_{p-1} - \mathbf{M}[v])^{-2}T[v])_u, \quad (3.4)$$

the u -entry of the column vector $(I_{p-1} - \mathbf{M}[v])^{-2}T[v]$.

Proof. We first give a “formal” argument and then justify its validity. The probability that when we take n steps from u , we never reach v and end up at some vertex w is $(\mathbf{M}[v]^n)_{uw}$ [why?]. The probability that once we reach w the next step is to v is $\mu(w, v)/d_w$. Hence by definition of expectation we have

$$H(u, v) = \sum_{w \neq v} \sum_{n \geq 0} (n+1) \frac{\mu(w, v)}{d_w} (\mathbf{M}[v]^n)_{uw}. \quad (3.5)$$

We claim that if x is a complex number satisfying $|x| < 1$, then

$$\sum_{n \geq 0} (n+1)x^n = (1-x)^{-2}. \quad (3.6)$$

This identity is a simple exercise in calculus. For instance, we can compute the coefficient of x^n in the product $(1-x)^2 \sum_{n \geq 0} (n+1)x^n$. We can also differentiate the familiar identity

$$\sum_{n \geq 0} x^n = \frac{1}{1-x}. \quad (3.7)$$

Another proof is obtained by expanding $(1-x)^{-2}$ by the binomial theorem for the exponent -2 . Convergence for $|x| < 1$ follows for example from the corresponding result for equation (3.7).

Let us “blindly” apply (3.6) to equation (3.5). We obtain

$$\begin{aligned} H(u, v) &= \sum_{w \neq v} ((I_{p-1} - \mathbf{M}[v])^{-2})_{ww} \frac{\mu(w, v)}{d_w} \\ &= ((I_{p-1} - \mathbf{M}[v])^{-2} T[v])_u, \end{aligned} \quad (3.8)$$

as claimed.

It remains to justify our derivation of equation (3.8). For an arbitrary real (or complex) $r \times r$ matrix \mathbf{B} , we can define $\sum_{n \geq 0} (n+1)\mathbf{B}^n$ entry-wise, that is, we set $\sum_{n \geq 0} (n+1)\mathbf{B}^n = \mathbf{C}$ if

$$\sum_{n \geq 0} (n+1)(\mathbf{B}^n)_{ij} = \mathbf{C}_{ij}$$

for all i and j indexing the rows and columns of \mathbf{B} and \mathbf{C} .

It is straightforward to verify by induction on m the identity

$$(I_r - \mathbf{B})^2 (I_r + 2\mathbf{B} + 3\mathbf{B}^2 + \cdots + m\mathbf{B}^{m-1}) = I_r - (m+1)\mathbf{B}^m + m\mathbf{B}^{m-1}. \quad (3.9)$$

Suppose that \mathbf{B} is diagonalizable and that all eigenvalues $\lambda_1, \dots, \lambda_r$ of \mathbf{B} satisfy $|\lambda_j| < 1$. Note that our proof of equation (1.1) extends to any diagonalizable matrix. (The matrix U need not be orthogonal, but this is irrelevant to the proof.) Hence

$$(\mathbf{B}^n)_{ij} = c_1 \lambda_1^n + \cdots + c_r \lambda_r^n,$$

where c_1, \dots, c_r are complex numbers (independent from n). Hence from equation (3.9) we see that the limit as $m \rightarrow \infty$ of the right-hand side approaches I_r . It follows [why?] that $\sum_{n \geq 0} (n+1)\mathbf{B}^n$ converges to $(I_r - \mathbf{B})^{-2}$.

NOTE. The above argument shows that $I_r - \mathbf{B}$ is indeed invertible. This fact is also an immediate consequence of the hypothesis that all eigenvalues of \mathbf{B} have absolute value less than one, since in particular there is no eigenvalue $\lambda = 1$.

From the discussion above, it remains to show that $\mathbf{M}[v]$ is diagonalizable, with all eigenvalues of absolute value less than one. The diagonalizability of $\mathbf{M}[v]$ is shown in exactly the same way as for \mathbf{M} in Theorem 3.2. (Thus we see also that $\mathbf{M}[v]$ has real eigenvalues, though we don't need this fact here.) It remains to show that the eigenvalues $\theta_1, \dots, \theta_{p-1}$ of $\mathbf{M}[v]$ satisfy $|\theta_j| < 1$. We would like to apply Theorem 3.3 to the matrix $\mathbf{M}[v]$, but this matrix might not be irreducible since the graph $G - v$ (defined by deleting from G the vertex v and all incident edges) need not be connected or may be just an isolated vertex. If $G - v$ has connected components H_1, \dots, H_m , then we can order the vertices of $G - v$ so that $\mathbf{M}[v]$ has the block structure

$$\mathbf{M}[v] = \begin{bmatrix} \mathbf{N}_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{N}_2 & \cdots & \mathbf{0} \\ & & \vdots & \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{N}_m \end{bmatrix},$$

where each \mathbf{N}_i is irreducible or is the 1×1 matrix $[0]$ (corresponding to H_i being an isolated vertex). The eigenvalues of $\mathbf{M}[v]$ are the eigenvalues of the \mathbf{N}_i 's.

We need to show that each eigenvalue of \mathbf{N}_i has absolute value less than one. If $\mathbf{N}_i = [0]$ then the only eigenvalue is 0, so we may assume that H_i is not an isolated vertex. Suppose that H_i has k vertices, so \mathbf{N}_i is a $k \times k$ matrix. Let ρ_i be the largest real eigenvalue of \mathbf{N}_i , so by Theorem 3.3 all eigenvalues λ of \mathbf{N}_i satisfy $|\lambda| \leq \rho_i$. Let $U = [u_1, \dots, u_k]$ be an eigenvector for ρ_i with positive entries (which exists by Theorem 3.3). We regard U as a column vector. Let V be the row vector of length k of all 1's. Consider the matrix product $V\mathbf{N}_iU$. On the one hand we have

$$V\mathbf{N}_iU = V(\rho_iU) = \rho_i(u_1 + \cdots + u_k). \quad (3.10)$$

On the other hand, if σ_j denotes the j th column sum of \mathbf{N}_i , then

$$V\mathbf{N}_iU = [\sigma_1, \dots, \sigma_k]U = \sigma_1u_1 + \cdots + \sigma_ku_k. \quad (3.11)$$

Now every σ_j satisfies $0 \leq \sigma_j \leq 1$, and at least one σ_h satisfies $\sigma_h < 1$ [why?]. Since each $u_j > 0$, it follows from equation (3.11) that $V\mathbf{N}_i U < u_1 + \cdots + u_k$. Comparing with equation (3.10) gives $\rho_i < 1$.

Since the eigenvalues of $\mathbf{M}[v]$ are just the eigenvalues of the \mathbf{N}_i 's, we see that all eigenvalues θ of $\mathbf{M}[v]$ satisfy $|\theta| < 1$. This completes the proof of Theorem 3.4. □

3.5 Example. Let G be the graph of Figure 3.1 with $v = v_4$. Then

$$\begin{aligned} \mathbf{M} &= \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & 0 & \frac{1}{3} \\ \frac{1}{4} & 0 & \frac{1}{4} & \frac{1}{2} \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{4} & \frac{1}{2} & \frac{1}{4} & 0 \end{bmatrix} \\ I_3 - \mathbf{M}[v] &= \begin{bmatrix} \frac{2}{3} & -\frac{1}{3} & 0 \\ -\frac{1}{4} & 1 & -\frac{1}{4} \\ 0 & -\frac{1}{2} & 1 \end{bmatrix} \\ (I_3 - \mathbf{M}[v])^{-2} &= \begin{bmatrix} \frac{55}{16} & \frac{13}{6} & \frac{17}{24} \\ \frac{13}{8} & \frac{7}{3} & \frac{11}{12} \\ \frac{17}{16} & \frac{11}{6} & \frac{13}{8} \end{bmatrix} \\ (I_3 - \mathbf{M}[v])^{-2} \begin{bmatrix} \frac{1}{3} \\ \frac{1}{2} \\ \frac{1}{2} \end{bmatrix} &= \begin{bmatrix} \frac{31}{12} \\ \frac{13}{6} \\ \frac{25}{12} \end{bmatrix}. \end{aligned}$$

Thus $H(v_1, v) = 31/12$, $H(v_2, v) = 13/6$, and $H(v_3, v) = 25/12$.

NOTE. The method used to prove that $\sum_{n \geq 0} (n+1)\mathbf{B}^n$ converges when all eigenvalues of \mathbf{B} have absolute value less than one can be extended, with a little more work (mostly concerned with non-diagonalizability), to show the following. Let $F(x) = \sum_{n \geq 0} a_n x^n$ be a power series with complex coefficients a_n . Let $\alpha > 0$ be such that $F(x)$ converges whenever $|x| < \alpha$. Let \mathbf{B} be a square matrix (over the complex numbers) whose eigenvalues λ all satisfy $|\lambda| < \alpha$. Then the matrix power series $\sum_{n \geq 0} a_n \mathbf{B}^n$ converges in the entry-wise sense described above.

Notes for Chapter 3




Figure 3.1: A graph for Example 3.5

Random walks on graphs is a vast subject, of which we have barely scratched the surface. Two typical questions considerably deeper than what we have considered are the following: how rapidly does a random walk approach the stationary distribution of Exercise 3.1? Assuming G is connected, what is the expected number of steps needed to visit every vertex? For a nice survey of random walks in graphs, see Lovász [73]. The topic of matrix power series is part of the subject of *matrix analysis*. For further information, see for instance Chapter 5 of the text by Horn and Johnson [59]. Our proof of Theorem 3.4 is somewhat “naive,” avoiding the development of the theory of matrix norms.

Chapter 4

The Sperner property

In this chapter we consider a surprising application of certain adjacency matrices to some problems in extremal set theory. An important role will also be played by finite groups in Chapter 5, which is a continuation of the present chapter. In general, extremal set theory is concerned with finding (or estimating) the most or least number of sets satisfying given set-theoretic or combinatorial conditions. For example, a typical easy problem in extremal set theory is the following: what is the most number of subsets of an n -element set with the property that any two of them intersect? (Can you solve this problem?) The problems to be considered here are most conveniently formulated in terms of *partially ordered sets*, or posets for short. Thus we begin with discussing some basic notions concerning posets.


4.1 Definition. A *poset* P is a finite set, also denoted P , together with a binary relation denoted \leq satisfying the following axioms:

- (P1) (reflexivity) $x \leq x$ for all $x \in P$.
- (P2) (antisymmetry) If $x \leq y$ and $y \leq x$, then $x = y$.
- (P3) (transitivity) If $x \leq y$ and $y \leq z$, then $x \leq z$.

One easy way to obtain a poset is the following. Let P be any collection of sets. If $x, y \in P$, then define $x \leq y$ in P if $x \subseteq y$ as sets. It is easy to see that this definition of \leq makes P into a poset. If P consists of *all* subsets of an n -element set S , then P is called a (finite) *boolean algebra* of rank n and is denoted by B_S . If $S = \{1, 2, \dots, n\}$, then we denote B_S simply by

B_n . Boolean algebras will play an important role throughout this chapter and the next.

There is a simple way to represent small posets pictorially. The *Hasse diagram* of a poset P is a planar drawing, with elements of P drawn as dots. If $x < y$ in P (i.e., $x \leq y$ and $x \neq y$), then y is drawn “above” x (i.e., with a larger vertical coordinate). An edge is drawn between x and y if y covers x , i.e., $x < y$ and no element z satisfies $x < z < y$. We then write $x \leq y$ or $y \geq x$. By the transitivity property (P3), all the relations of a finite poset are determined by the cover relations, so the Hasse diagram determines P . (This is not true for infinite posets; for instance, the real numbers \mathbb{R} with their usual order is a poset with no cover relations.) The Hasse diagram of the boolean algebra B_3 looks like



We say that two posets P and Q are *isomorphic* if there is a bijection (one-to-one and onto function) $\varphi: P \rightarrow Q$ such that $x \leq y$ in P if and only if $\varphi(x) \leq \varphi(y)$ in Q . Thus one can think that two posets are isomorphic if they differ only in the names of their elements. This is exactly analogous to the notion of isomorphism of groups, rings, etc. It is an instructive exercise (see Exercise 4.1) to draw Hasse diagrams of the one poset of *order* (number of elements) one (up to isomorphism), the two posets of order two, the five posets of order three, and the sixteen posets of order four. More ambitious readers can try the 63 posets of order five, the 318 of order six, the 2045 of order seven, the 16999 of order eight, the 183231 of order nine, the 2567284 of order ten, the 46749427 of order eleven, the 1104891746 of order twelve, the 33823827452 of order thirteen, the 1338193159771 of order fourteen, the 68275077901156 of order fifteen, and the 4483130665195087 of order sixteen. Beyond this the number is not currently known.

A *chain* C in a poset is a totally ordered subset of P , i.e., if $x, y \in C$ then

either $x \leq y$ or $y \leq x$ in P . A finite chain is said to have *length* n if it has $n + 1$ elements. Such a chain thus has the form $x_0 < x_1 < \cdots < x_n$. We say that a finite poset is *graded of rank* n if every maximal chain has length n . (A chain is *maximal* if it's contained in no larger chain.) For instance, the boolean algebra B_n is graded of rank n [why?]. A chain $y_0 < y_1 < \cdots < y_j$ is said to be *saturated* if each y_{i+1} covers y_i . Such a chain need not be maximal since there can be elements of P less than y_0 or greater than y_j . If P is graded of rank n and $x \in P$, then we say that x has *rank* j , denoted $\rho(x) = j$, if the largest saturated chain of P with top element x has length j . Thus [why?] if we let $P_j = \{x \in P: \rho(x) = j\}$, then P is a *disjoint union* $P = P_0 \cup P_1 \cup \cdots \cup P_n$, and every maximal chain of P has the form $x_0 < x_1 < \cdots < x_n$ where $\rho(x_j) = j$. We call P_i the *i*th *level* of P . We write $p_j = \#P_j$, the number of elements of P of rank j . For example, if $P = B_n$ then $\rho(x) = |x|$ (the cardinality of x as a set) and

$$p_j = \#\{x \subseteq \{1, 2, \dots, n\}: |x| = j\} = \binom{n}{j}.$$

(Note that we use both $|S|$ and $\#S$ for the cardinality of a finite set S .) If a graded poset P of rank n has p_i elements of rank i , then define the *rank-generating function*

$$F(P, q) = \sum_{i=0}^n p_i q^i = \sum_{x \in P} q^{\rho(x)}.$$

For instance, $F(B_n, q) = (1 + q)^n$ [why?].

We say that a graded poset P of rank n (always assumed to be finite) is *rank-symmetric* if $p_i = p_{n-i}$ for $0 \leq i \leq n$, and *rank-unimodal* if $p_0 \leq p_1 \leq \cdots \leq p_j \geq p_{j+1} \geq p_{j+2} \geq \cdots \geq p_n$ for some $0 \leq j \leq n$. If P is both rank-symmetric and rank-unimodal, then we clearly have

$$p_0 \leq p_1 \leq \cdots \leq p_m \geq p_{m+1} \geq \cdots \geq p_n, \text{ if } n = 2m$$

$$p_0 \leq p_1 \leq \cdots \leq p_m = p_{m+1} \geq p_{m+2} \geq \cdots \geq p_n, \text{ if } n = 2m + 1.$$

We also say that the sequence p_0, p_1, \dots, p_n itself or the polynomial $F(q) = p_0 + p_1 q + \cdots + p_n q^n$ is *symmetric* or *unimodal*, as the case may be. For instance, B_n is rank-symmetric and rank-unimodal, since it is well-known (and easy to prove) that the sequence $\binom{n}{0}, \binom{n}{1}, \dots, \binom{n}{n}$ (the n th row of Pascal's triangle) is symmetric and unimodal. Thus the polynomial $(1 + q)^n$ is symmetric and unimodal.

A few more definitions, and then finally some results! An *antichain* in a poset P is a subset A of P for which no two elements are comparable, i.e., we can never have $x, y \in A$ and $x < y$. For instance, in a graded poset P the “levels” P_j are antichains [why?]. We will be concerned with the problem of finding the largest antichain in a poset. Consider for instance the boolean algebra B_n . The problem of finding the largest antichain in B_n is clearly equivalent to the following problem in extremal set theory: find the largest collection of subsets of an n -element set such that no element of the collection contains another. A good guess would be to take all the subsets of cardinality $\lfloor n/2 \rfloor$ (where $\lfloor x \rfloor$ denotes the greatest integer $\leq x$), giving a total of $\binom{n}{\lfloor n/2 \rfloor}$ sets in all. But how can we actually prove there is no larger collection? Such a proof was first given by Emanuel Sperner in 1927 and is known as *Sperner’s theorem*. We will give three proofs of Sperner’s theorem in this chapter: one proof uses linear algebra and will be applied to certain other situations; the second proof is an elegant combinatorial argument due to David Lubell in 1966; while the third proof is another combinatorial argument closely related to the linear algebra proof. We present the last two proofs for their “cultural value.” Our extension of Sperner’s theorem to certain other situations will involve the following crucial definition.


4.2 Definition. Let P be a graded poset of rank n . We say that P has the *Sperner property* or is a *Sperner poset* if

$$\max\{\#A : A \text{ is an antichain of } P\} = \max\{\#P_i : 0 \leq i \leq n\}.$$

In other words, no antichain is larger than the largest level P_i .

Thus Sperner’s theorem is equivalent to saying that B_n has the Sperner property. Note that if P has the Sperner property then there may still be antichains of maximum cardinality other than the biggest P_i ; there just can’t be any bigger antichains.

4.3 Example. A simple example of a graded poset that fails to satisfy the Sperner property is the following:



We now will discuss a simple combinatorial condition which guarantees that certain graded posets P are Sperner. We define an *order-matching* from P_i to P_{i+1} to be a *one-to-one* function $\mu: P_i \rightarrow P_{i+1}$ satisfying $x < \mu(x)$ for all $x \in P_i$. Clearly if such an order-matching exists then $p_i \leq p_{i+1}$ (since μ is one-to-one). Easy examples (such as the diagram above) show that the converse is false, i.e., if $p_i \leq p_{i+1}$ then there need not exist an order-matching from P_i to P_{i+1} . We similarly define an order-matching from P_i to P_{i-1} to be a one-to-one function $\mu: P_i \rightarrow P_{i-1}$ satisfying $\mu(x) < x$ for all $x \in P_i$.

4.4 Proposition. *Let P be a graded poset of rank n . Suppose there exists an integer $0 \leq j \leq n$ and order-matchings*

$$P_0 \rightarrow P_1 \rightarrow P_2 \rightarrow \cdots \rightarrow P_j \leftarrow P_{j+1} \leftarrow P_{j+2} \leftarrow \cdots \leftarrow P_n. \quad (4.1)$$

Then P is rank-unimodal and Sperner.

Proof. Since order-matchings are one-to-one it is clear that

$$p_0 \leq p_1 \leq \cdots \leq p_j \geq p_{j+1} \geq p_{j+2} \geq \cdots \geq p_n.$$

Hence P is rank-unimodal.

Define a graph G as follows. The vertices of G are the elements of P . Two vertices x, y are connected by an edge if one of the order-matchings μ in the statement of the proposition satisfies $\mu(x) = y$. (Thus G is a subgraph of the Hasse diagram of P .) Drawing a picture will convince you that G consists of a disjoint union of paths, including single-vertex paths not involved in any of the order-matchings. The vertices of each of these paths form a chain in P . Thus we have partitioned the elements of P into disjoint chains. Since P is rank-unimodal with biggest level P_j , all of these chains must pass through P_j [why?]. Thus the number of chains is exactly p_j . Any antichain A can intersect each of these chains at most once, so the cardinality $|A|$ of A cannot exceed the number of chains, i.e., $|A| \leq p_j$. Hence by definition P is Sperner. \square

It is now finally time to bring some linear algebra into the picture. For any (finite) set S , we let $\mathbb{R}S$ denote the real vector space consisting of all formal linear combinations (with real coefficients) of elements of S . Thus S is a basis for $\mathbb{R}S$, and in fact we could have simply defined $\mathbb{R}S$ to be the real vector space with basis S . The next lemma relates the combinatorics we have just discussed to linear algebra and will allow us to prove that certain posets are Sperner by the use of linear algebra (combined with some finite group theory).

4.5 Lemma. *Suppose there exists a linear transformation $U: \mathbb{R}P_i \rightarrow \mathbb{R}P_{i+1}$ (U stands for “up”) satisfying:*

- U is one-to-one.
- For all $x \in P_i$, $U(x)$ is a linear combination of elements $y \in P_{i+1}$ satisfying $x < y$. (We then call U an order-raising operator.)

Then there exists an order-matching $\mu: P_i \rightarrow P_{i+1}$.

Similarly, suppose there exists a linear transformation $U: \mathbb{R}P_i \rightarrow \mathbb{R}P_{i+1}$ satisfying:

- U is onto.
- U is an order-raising operator.

Then there exists an order-matching $\mu: P_{i+1} \rightarrow P_i$.

Proof. Suppose $U: \mathbb{R}P_i \rightarrow \mathbb{R}P_{i+1}$ is a one-to-one order-raising operator. Let $[U]$ denote the matrix of U with respect to the bases P_i of $\mathbb{R}P_i$ and P_{i+1} of $\mathbb{R}P_{i+1}$. Thus the rows of $[U]$ are indexed by the elements $y_1, \dots, y_{p_{i+1}}$ of P_{i+1} (in some order), and the columns by the elements x_1, \dots, x_{p_i} of P_i . Since U is one-to-one, the rank of $[U]$ is equal to p_i (the number of columns). Since the row rank of a matrix equals its column rank, $[U]$ must have p_i linearly independent rows. Say we have labelled the elements of P_{i+1} so that the first p_i rows of $[U]$ are linearly independent.

Let $A = (a_{ij})$ be the $p_i \times p_i$ matrix whose rows are the first p_i rows of $[U]$. (Thus A is a square submatrix of $[U]$.) Since the rows of A are linearly independent, we have

$$\det(A) = \sum \pm a_{1\pi(1)} \cdots a_{p_i\pi(p_i)} \neq 0,$$

where the sum is over all permutations π of $1, \dots, p_i$. Thus some term $\pm a_{1\pi(1)} \cdots a_{p_i\pi(p_i)}$ of the above sum is nonzero. Since U is order-raising, this means that [why?] $y_k > x_{\pi(k)}$ for $1 \leq k \leq p_i$. Hence the map $\mu: P_i \rightarrow P_{i+1}$ defined by $\mu(x_k) = y_{\pi^{-1}(k)}$ is an order-matching, as desired.

The case when U is onto rather than one-to-one is proved by a completely analogous argument. It can also be deduced from the one-to-one case by considering the transpose of the matrix $[U]$. \square

NOTE. Although it does not really help in understanding the theory, it is interesting to regard a one-to-one order-raising operator as a “quantum order-matching.” Rather than choosing a single element $y = \mu(x)$ that is matched with $x \in P_i$, we choose all possible elements $y \in P_{i+1}$ satisfying $y > x$ at the same time. If $U(x) = \sum_{y>x} c_y y$ (where $c_y \in \mathbb{R}$), then we are choosing y with “weight” c_y . As explained in the proof of Lemma 4.5 above, we “break the symmetry” and obtain a single matched element $\mu(x)$ by choosing some nonvanishing term in the expansion of a determinant.

We now want to apply Proposition 4.4 and Lemma 4.5 to the boolean algebra B_n . For each $0 \leq i < n$, we need to define a linear transformation $U_i: \mathbb{R}(B_n)_i \rightarrow \mathbb{R}(B_n)_{i+1}$, and then prove it has the desired properties. We simply define U_i to be the simplest possible order-raising operator, namely, for $x \in (B_n)_i$, let

$$U_i(x) = \sum_{\substack{y \in (B_n)_{i+1} \\ y > x}} y. \quad (4.2)$$

Note that since $(B_n)_i$ is a basis for $\mathbb{R}(B_n)_i$, equation (4.2) does indeed define a unique linear transformation $U_i: \mathbb{R}(B_n)_i \rightarrow \mathbb{R}(B_n)_{i+1}$. By definition U_i is order-raising; we want to show that U_i is one-to-one for $i < n/2$ and onto for $i \geq n/2$. There are several ways to show this using only elementary linear algebra; we will give what is perhaps the simplest proof, though it is quite tricky. The idea is to introduce “dual” or “adjoint” operators $D_i: \mathbb{R}(B_n)_i \rightarrow \mathbb{R}(B_n)_{i-1}$ to the U_i ’s (D stands for “down”), defined by

$$D_i(y) = \sum_{\substack{x \in (B_n)_{i-1} \\ x < y}} x, \quad (4.3)$$

for all $y \in (B_n)_i$. Let $[U_i]$ denote the matrix of U_i with respect to the bases $(B_n)_i$ and $(B_n)_{i+1}$, and similarly let $[D_i]$ denote the matrix of D_i with respect to the bases $(B_n)_i$ and $(B_n)_{i-1}$. A key observation which we will use later is that

$$[D_{i+1}] = [U_i]^t, \quad (4.4)$$

i.e., the matrix $[D_{i+1}]$ is the transpose of the matrix $[U_i]$ [why?]. Now let $I_i: \mathbb{R}(B_n)_i \rightarrow \mathbb{R}(B_n)_i$ denote the identity transformation on $\mathbb{R}(B_n)_i$, i.e., $I_i(u) = u$ for all $u \in \mathbb{R}(B_n)_i$. The next lemma states (in linear algebraic terms) the fundamental combinatorial property of B_n which we need. For this lemma set $U_n = 0$ and $D_0 = 0$ (the 0 linear transformation between the appropriate vector spaces).

4.6 Lemma. *Let $0 \leq i \leq n$. Then*

$$D_{i+1}U_i - U_{i-1}D_i = (n - 2i)I_i. \quad (4.5)$$

(Linear transformations are multiplied right-to-left, so $AB(u) = A(B(u))$.)

Proof. Let $x \in (B_n)_i$. We need to show that if we apply the left-hand side of equation (4.5) to x , then we obtain $(n - 2i)x$. We have

$$\begin{aligned} D_{i+1}U_i(x) &= D_{i+1} \left(\sum_{\substack{|y|=i+1 \\ x \subset y}} y \right) \\ &= \sum_{\substack{|y|=i+1 \\ x \subset y}} \sum_{\substack{|z|=i \\ z \subset y}} z. \end{aligned}$$

If $x, z \in (B_n)_i$ satisfy $|x \cap z| < i - 1$, then there is no $y \in (B_n)_{i+1}$ such that $x \subset y$ and $z \subset y$. Hence the coefficient of z in $D_{i+1}U_i(x)$ when it is expanded in terms of the basis $(B_n)_i$ is 0. If $|x \cap z| = i - 1$, then there is one such y , namely, $y = x \cup z$. Finally if $x = z$ then y can be any element of $(B_n)_{i+1}$ containing x , and there are $n - i$ such y in all. It follows that

$$D_{i+1}U_i(x) = (n - i)x + \sum_{\substack{|z|=i \\ |x \cap z|=i-1}} z. \quad (4.6)$$

By exactly analogous reasoning (which the reader should check), we have for $x \in (B_n)_i$ that

$$U_{i-1}D_i(x) = ix + \sum_{\substack{|z|=i \\ |x \cap z|=i-1}} z. \quad (4.7)$$

Subtracting (4.7) from (4.6) yields $(D_{i+1}U_i - U_{i-1}D_i)(x) = (n - 2i)x$, as desired. \square

4.7 Theorem. *The operator U_i defined above is one-to-one if $i < n/2$ and is onto if $i \geq n/2$.*

Proof. Recall that $[D_i] = [U_{i-1}]^t$. From linear algebra we know that a (rectangular) matrix times its transpose is *positive semidefinite* (or just *semidefinite* for short) and hence has nonnegative (real) eigenvalues. By Lemma 4.6 we have

$$D_{i+1}U_i = U_{i-1}D_i + (n - 2i)I_i.$$

Thus the eigenvalues of $D_{i+1}U_i$ are obtained from the eigenvalues of $U_{i-1}D_i$ by adding $n - 2i$. Since we are assuming that $n - 2i > 0$, it follows that the eigenvalues of $D_{i+1}U_i$ are strictly positive. Hence $D_{i+1}U_i$ is invertible (since it has no 0 eigenvalues). But this implies that U_i is one-to-one [why?], as desired.

The case $i \geq n/2$ is done by a “dual” argument (or in fact can be deduced directly from the $i < n/2$ case by using the fact that the poset B_n is “self-dual,” though we will not go into this). Namely, from the fact that

$$U_i D_{i+1} = D_{i+2} U_{i+1} + (2i + 2 - n) I_{i+1}$$

we get that $U_i D_{i+1}$ is invertible, so now U_i is onto, completing the proof. \square

Combining Proposition 4.4, Lemma 4.5, and Theorem 4.7, we obtain the famous theorem of Sperner.

4.8 Corollary. *The boolean algebra B_n has the Sperner property.*

It is natural to ask whether there is a less indirect proof of Corollary 4.8. In fact, several nice proofs are known; we first give one due to David Lubell, mentioned before Definition 4.2.

Lubell’s proof of Sperner’s theorem. First we count the total number of maximal chains $\emptyset = x_0 < x_1 < \cdots < x_n = \{1, \dots, n\}$ in B_n . There are n choices for x_1 , then $n - 1$ choices for x_2 , etc., so there are $n!$ maximal chains in all. Next we count the number of maximal chains $x_0 < x_1 < \cdots < x_i = x < \cdots < x_n$ which contain a given element x of rank i . There are i choices for x_1 , then $i - 1$ choices for x_2 , up to one choice for x_i . Similarly there are $n - i$ choices for x_{i+1} , then $n - i - 1$ choices for x_{i+2} , etc., up to one choice for x_n . Hence the number of maximal chains containing x is $i!(n - i)!$.

Now let A be an antichain. If $x \in A$, then let C_x be the set of maximal chains of B_n which contain x . Since A is an antichain, the sets C_x , $x \in A$ are pairwise disjoint. Hence

$$\begin{aligned} \left| \bigcup_{x \in A} C_x \right| &= \sum_{x \in A} |C_x| \\ &= \sum_{x \in A} (\rho(x))!(n - \rho(x))! \end{aligned}$$

Since the total number of maximal chains in the C_x 's cannot exceed the total number $n!$ of maximal chains in B_n , we have

$$\sum_{x \in A} (\rho(x))!(n - \rho(x))! \leq n!.$$

Divide both sides by $n!$ to obtain

$$\sum_{x \in A} \frac{1}{\binom{n}{\rho(x)}} \leq 1.$$

Since $\binom{n}{i}$ is maximized when $i = \lfloor n/2 \rfloor$, we have

$$\frac{1}{\binom{n}{\lfloor n/2 \rfloor}} \leq \frac{1}{\binom{n}{\rho(x)}},$$

for all $x \in A$ (or all $x \in B_n$). Thus

$$\sum_{x \in A} \frac{1}{\binom{n}{\lfloor n/2 \rfloor}} \leq 1,$$

or equivalently,

$$|A| \leq \binom{n}{\lfloor n/2 \rfloor}.$$

Since $\binom{n}{\lfloor n/2 \rfloor}$ is the size of the largest level of B_n , it follows that B_n is Sperner. \square

There is another nice way to show directly that B_n is Sperner, namely, by constructing an explicit order-matching $\mu: (B_n)_i \rightarrow (B_n)_{i+1}$ when $i < n/2$. We will define μ by giving an example. Let $n = 21$, $i = 9$, and $S = \{3, 4, 5, 8, 12, 13, 17, 19, 20\}$. We want to define $\mu(S)$. Let $(a_1, a_2, \dots, a_{21})$ be a sequence of ± 1 's, where $a_i = 1$ if $i \in S$, and $a_i = -1$ if $i \notin S$. For the set S above we get the sequence (writing $-$ for -1)

$$- - 1 1 1 - - 1 - - - 1 1 - - - 1 - 1 1 - .$$

Replace any two consecutive terms $1 -$ with 00 :

$$- - 1 1 0 0 - 0 0 - - 1 0 0 - - 0 0 1 0 0 .$$

Ignore the 0's and replace any two consecutive terms $1 -$ with 00 :

$$- - 1 0 0 0 0 0 - - 0 0 0 0 - 0 0 1 0 0 .$$

Continue:

$$- - 00000000 - 0000 - 00100.$$

At this stage no further replacement is possible. The nonzero terms consist of a sequence of $-$'s followed by a sequence of 1 's. There is at least one $-$ since $i < n/2$. Let k be the position (coordinate) of the last $-$; here $k = 16$. Define $\mu(S) = S \cup \{k\} = S \cup \{16\}$. The reader can check that this procedure gives an order-matching. In particular, why is μ injective (one-to-one), i.e., why can we recover S from $\mu(S)$?

It can be checked that if we glue together the order-matchings $(B_n)_i \rightarrow (B_n)_{i+1}$ for $i < n/2$ just defined, along with an obvious dual construction $(B_n)_i \rightarrow (B_n)_{i-1}$ for $i > n/2$ then we obtain more than just a partition of B_n into saturated chains passing through the middle level (n even) or middle two levels (n odd), as in the proof of Proposition 4.4. We in fact have the additional property that these chains are all *symmetric*, i.e., they begin at some level $i \leq n/2$ and end at level $n - i$. Such a decomposition of a rank-symmetric, rank-unimodal graded poset P into saturated chains is called a *symmetric chain decomposition*. A symmetric chain decomposition implies that for any $j \geq 1$, the largest size of a union of j antichains is equal to the largest size of a union of j levels of P (Exercise 4.6). (The Sperner property corresponds to the case $j = 1$). It can be a challenging problem to decide whether certain posets have a symmetric chain decomposition (e.g., Exercises 5.5(b), 5.6 and 6.6), though we will not discuss this topic further here.

In view of the above elegant proof of Lubell and the explicit description of an order-matching $\mu : (B_n)_i \rightarrow (B_n)_{i+1}$, the reader may be wondering what was the point of giving a rather complicated and indirect proof using linear algebra. Admittedly, if all we could obtain from the linear algebra machinery we have developed was just another proof of Sperner's theorem, then it would have been hardly worth the effort. But in the next chapter we will show how Theorem 4.7, when combined with a little finite group theory, can be used to obtain many interesting combinatorial results for which simple, direct proofs are not known.

Notes for Chapter 4

For further information on combinatorial aspects of partially ordered sets in general, see Caspard-Leclerc-Monjardet [21], P. Fishburn [34], R. Stanley

[109, Ch. 3], and W. Trotter [115]. Sperner's theorem (Corollary 4.8) was first proved by E. Sperner [102]. The elegant proof of Lubell appears in [74]. A general reference on the Sperner property is the book by K. Engel [33]. For more general results on the combinatorics of finite sets, see I. Anderson [2]. The linear algebraic approach to the Sperner property discussed here is due independently to M. Pouzet [88] (further developed by Pouzet and Rosenberg [89]) and R. Stanley [103][105]. For further information on explicit order matchings, symmetric chain decompositions, etc., see the text [2] of Anderson mentioned above.

Chapter 5

Group actions on boolean algebras

Let us begin by reviewing some facts from group theory. Suppose that X is an n -element set and that G is a group. We say that G *acts on* the set X if for every element π of G we associate a permutation (also denoted π) of X , such that for all $x \in X$ and $\pi, \sigma \in G$ we have

$$\pi(\sigma(x)) = (\pi\sigma)(x).$$

Thus [why?] an action of G on X is the same as a homomorphism $\varphi: G \rightarrow \mathfrak{S}_X$, where \mathfrak{S}_X denotes the symmetric group of all permutations of X . We sometimes write $\pi \cdot x$ instead of $\pi(x)$.

5.1 Example. (a) Let the real number α act on the xy -plane by rotation counterclockwise around the origin by an angle of α radians. It is easy to check that this defines an action of the group \mathbb{R} of real numbers (under addition) on the xy -plane. The kernel of this action, i.e., the kernel of the homomorphism $\varphi: \mathbb{R} \rightarrow \mathfrak{S}_{\mathbb{R}^2}$, is the cyclic subgroup of \mathbb{R} generated by 2π .

(b) Now let $\alpha \in \mathbb{R}$ act by translation by a distance α to the right, i.e., adding $(\alpha, 0)$. This yields a completely different action of \mathbb{R} on the xy -plane. This time the action is *faithful*, i.e., the kernel is the trivial subgroup $\{0\}$.

(c) Let $X = \{a, b, c, d\}$ and $G = \mathbb{Z}_2 \times \mathbb{Z}_2 = \{(0, 0), (0, 1), (1, 0), (1, 1)\}$. Let G act as follows:

$$(0, 1) \cdot a = b, \quad (0, 1) \cdot b = a, \quad (0, 1) \cdot c = c, \quad (0, 1) \cdot d = d$$

$$(1, 0) \cdot a = a, \quad (1, 0) \cdot b = b, \quad (1, 0) \cdot c = d, \quad (1, 0) \cdot d = c.$$

The reader should check that this does indeed define an action. In particular, since $(1, 0)$ and $(0, 1)$ generate G , we don't need to define the action of $(0, 0)$ and $(1, 1)$ — they are uniquely determined.

(d) Let X and G be as in (c), but now define the action by

$$(0, 1) \cdot a = b, \quad (0, 1) \cdot b = a, \quad (0, 1) \cdot c = d, \quad (0, 1) \cdot d = c$$

$$(1, 0) \cdot a = c, \quad (1, 0) \cdot b = d, \quad (1, 0) \cdot c = a, \quad (1, 0) \cdot d = b.$$

Again one can check that we have an action of $\mathbb{Z}_2 \times \mathbb{Z}_2$ on $\{a, b, c, d\}$. The two actions of $G = \mathbb{Z}_2 \times \mathbb{Z}_2$ that we have just defined are quite different; for instance, in the first action we have some elements of X fixed by some nonidentity element of G (such as $(0, 1) \cdot c = c$), while the second action fails to have this property. See also Example 5.2(c,d) below for another fundamental way in which the two actions differ.

Recall what is meant by an *orbit* of the action of a group G on a set X . Namely, we say that two elements x, y of X are *G-equivalent* if $\pi(x) = y$ for some $\pi \in G$. The relation of G -equivalence is an equivalence relation, and the equivalence classes are called orbits. Thus x and y are in the same orbit if $\pi(x) = y$ for some $\pi \in G$. The orbits form a *partition* of X , i.e., they are pairwise-disjoint, nonempty subsets of X whose union is X . The orbit containing x is denoted Gx ; this is sensible notation since Gx consists of all elements $\pi(x)$ where $\pi \in G$. Thus $Gx = Gy$ if and only if x and y are G -equivalent (i.e., in the same G -orbit). The set of all G -orbits is denoted X/G .

5.2 Example. (a) In Example 5.1(a), the orbits are circles with center $(0, 0)$, including the degenerate circle whose only point is $(0, 0)$.

(b) In Example 5.1(b), the orbits are horizontal lines. Note that although in (a) and (b) the same group G acts on the same set X , the orbits are different.

(c) In Example 5.1(c), the orbits are $\{a, b\}$ and $\{c, d\}$.

(d) In Example 5.1(d), there is only one orbit $\{a, b, c, d\}$. Again we have a situation in which a group G acts on a set X in two different ways, with different orbits.

We wish to consider the situation where $X = B_n$, the boolean algebra of rank n (so $|B_n| = 2^n$). We begin by defining an *automorphism* of a poset P to be an isomorphism $\varphi: P \rightarrow P$. (This definition is exactly analogous

to the definition of an automorphism of a group, ring, etc.) The set of all automorphisms of P forms a group, denoted $\text{Aut}(P)$ and called the *automorphism group* of P , under the operation of composition of functions (just as is the case for groups, rings, etc.)

Now consider the case $P = B_n$. Any permutation π of $\{1, \dots, n\}$ acts on B_n as follows: if $x = \{i_1, i_2, \dots, i_k\} \in B_n$, then

$$\pi(x) = \{\pi(i_1), \pi(i_2), \dots, \pi(i_k)\}. \quad (5.1)$$


This action of π on B_n is an automorphism [why?]; in particular, if $|x| = i$, then also $|\pi(x)| = i$. Equation (5.1) defines an action of the symmetric group \mathfrak{S}_n of all permutations of $\{1, \dots, n\}$ on B_n [why?]. (In fact, it is not hard to show that *every* automorphism of B_n is of the form (5.1) for $\pi \in \mathfrak{S}_n$.) In particular, any subgroup G of \mathfrak{S}_n acts on B_n *via* (5.1) (where we restrict π to belong to G). In what follows this action is always meant.

5.3 Example. Let $n = 3$, and let G be the subgroup of \mathfrak{S}_3 with elements ι and $(1, 2)$. Here ι denotes the identity permutation, and (using disjoint cycle notation) $(1, 2)$ denotes the permutation which interchanges 1 and 2, and fixes 3. There are six orbits of G (acting on B_3). Writing e.g. 13 as short for $\{1, 3\}$, the six orbits are $\{\emptyset\}$, $\{1, 2\}$, $\{3\}$, $\{12\}$, $\{13, 23\}$, and $\{123\}$.


We now define the class of posets which will be of interest to us here. Later we will give some special cases of particular interest.

Let G be a subgroup of \mathfrak{S}_n . Define the *quotient poset* B_n/G as follows. The elements of B_n/G are the orbits of G . If \mathfrak{o} and \mathfrak{o}' are two orbits, then define $\mathfrak{o} \leq \mathfrak{o}'$ in B_n/G if there exist $x \in \mathfrak{o}$ and $y \in \mathfrak{o}'$ such that $x \leq y$ in B_n . It's easy to check that this relation \leq is indeed a partial order.

5.4 Example. (a) Let $n = 3$ and G be the group of order two generated by the cycle $(1, 2)$, as in Example 5.3. Then the Hasse diagram of B_3/G is shown below, where each element (orbit) is labeled by one of its elements.



(b) Let $n = 5$ and G be the group of order five generated by the cycle $(1, 2, 3, 4, 5)$. Then B_5/G has Hasse diagram



One simple property of a quotient poset B_n/G is the following.

5.5 Proposition. *The quotient poset B_n/G defined above is graded of rank n and rank-symmetric.*

Proof. We leave as an exercise the easy proof that B_n/G is graded of rank n , and that the rank of an element \mathfrak{o} of B_n/G is just the rank in B_n of any of the elements $x \in \mathfrak{o}$. Thus the number of elements $p_i(B_n/G)$ of rank i is

equal to the number of orbits $\mathfrak{o} \in (B_n)_i/G$. If $x \in B_n$, then let \bar{x} denote the set-theoretic complement of x , i.e.,

$$\bar{x} = \{1, \dots, n\} - x = \{1 \leq i \leq n: i \notin x\}.$$

Then $\{x_1, \dots, x_j\}$ is an orbit of i -element subsets of $\{1, \dots, n\}$ if and only if $\{\bar{x}_1, \dots, \bar{x}_j\}$ is an orbit of $(n-i)$ -element subsets [why?]. Hence $|(B_n)_i/G| = |(B_n)_{n-i}/G|$, so B_n/G is rank-symmetric. \square

Let $\pi \in \mathfrak{S}_n$. We associate with π a linear transformation (still denoted π) $\pi: \mathbb{R}(B_n)_i \rightarrow \mathbb{R}(B_n)_i$ by the rule

$$\pi \left(\sum_{x \in (B_n)_i} c_x x \right) = \sum_{x \in (B_n)_i} c_x \pi(x),$$

where each c_x is a real number. This defines an action of \mathfrak{S}_n , or of any subgroup G of \mathfrak{S}_n , on the vector space $\mathbb{R}(B_n)_i$. The matrix of π with respect to the basis $(B_n)_i$ is just a *permutation matrix*, i.e., a matrix with one 1 in every row and column, and 0's elsewhere. We will be interested in elements of $\mathbb{R}(B_n)_i$ which are fixed by every element of a subgroup G of \mathfrak{S}_n . The set of all such elements is denoted $\mathbb{R}(B_n)_i^G$, so

$$\mathbb{R}(B_n)_i^G = \{v \in \mathbb{R}(B_n)_i: \pi(v) = v \text{ for all } \pi \in G\}.$$

5.6 Lemma. *A basis for $\mathbb{R}(B_n)_i^G$ consists of the elements*

$$v_{\mathfrak{o}} := \sum_{x \in \mathfrak{o}} x,$$

where $\mathfrak{o} \in (B_n)_i/G$, the set of G -orbits for the action of G on $(B_n)_i$.

Proof. First note that if \mathfrak{o} is an orbit and $x \in \mathfrak{o}$, then by definition of orbit we have $\pi(x) \in \mathfrak{o}$ for all $\pi \in G$ (or all $\pi \in \mathfrak{S}_n$). Since π permutes the elements of $(B_n)_i$, it follows that π permutes the elements of \mathfrak{o} . Thus $\pi(v_{\mathfrak{o}}) = v_{\mathfrak{o}}$, so $v_{\mathfrak{o}} \in \mathbb{R}(B_n)_i^G$. It is clear that the $v_{\mathfrak{o}}$'s are linearly independent since any $x \in (B_n)_i$ appears with nonzero coefficient in exactly one $v_{\mathfrak{o}}$.

It remains to show that the $v_{\mathfrak{o}}$'s span $\mathbb{R}(B_n)_i^G$, i.e., any $v = \sum_{x \in (B_n)_i} c_x x \in \mathbb{R}(B_n)_i^G$ can be written as a linear combination of $v_{\mathfrak{o}}$'s. Given $x \in (B_n)_i$, let $G_x = \{\pi \in G: \pi(x) = x\}$, the *stabilizer* of x . We leave as an easy exercise

the standard fact that $\pi(x) = \sigma(x)$ (where $\pi, \sigma \in G$) if and only if π and σ belong to the same left coset of G_x , i.e., $\pi G_x = \sigma G_x$. It follows that in the multiset of elements $\pi(x)$, where π ranges over all elements of G and x is fixed, every element y in the orbit Gx appears $\#G_x$ times, and no other elements appear. In other words,

$$\sum_{\pi \in G} \pi(x) = |G_x| \cdot v_{Gx}.$$

(Do not confuse the orbit Gx with the subgroup G_x !) Now apply π to v and sum on all $\pi \in G$. Since $\pi(v) = v$ (because $v \in \mathbb{R}(B_n)_i^G$), we get

$$\begin{aligned} |G| \cdot v &= \sum_{\pi \in G} \pi(v) \\ &= \sum_{\pi \in G} \left(\sum_{x \in (B_n)_i} c_x \pi(x) \right) \\ &= \sum_{x \in (B_n)_i} c_x \left(\sum_{\pi \in G} \pi(x) \right) \\ &= \sum_{x \in (B_n)_i} c_x \cdot (\#G_x) \cdot v_{Gx}. \end{aligned}$$

Dividing by $|G|$ expresses v as a linear combination of the elements v_{Gx} (or v_o), as desired. \square

Now let us consider the effect of applying the order-raising operator U_i to an element v of $\mathbb{R}(B_n)_i^G$.

5.7 Lemma. *If $v \in \mathbb{R}(B_n)_i^G$, then $U_i(v) \in \mathbb{R}(B_n)_{i+1}^G$.*

Proof. Note that since $\pi \in G$ is an automorphism of B_n , we have $x < y$ in B_n if and only if $\pi(x) < \pi(y)$ in B_n . It follows [why?] that if $x \in (B_n)_i$ then

$$U_i(\pi(x)) = \pi(U_i(x)).$$

Since U_i and π are linear transformations, it follows by linearity that $U_i\pi(u) = \pi U_i(u)$ for all $u \in \mathbb{R}(B_n)_i$. In other words, $U_i\pi = \pi U_i$. Then

$$\begin{aligned} \pi(U_i(v)) &= U_i(\pi(v)) \\ &= U_i(v), \end{aligned}$$

so $U_i(v) \in \mathbb{R}(B_n)_{i+1}^G$, as desired. \square

We come to the main result of this chapter, and indeed our main result on the Sperner property.

5.8 Theorem. *Let G be a subgroup of \mathfrak{S}_n . Then the quotient poset B_n/G is graded of rank n , rank-symmetric, rank-unimodal, and Sperner.*

Proof. Let $P = B_n/G$. We have already seen in Proposition 5.5 that P is graded of rank n and rank-symmetric. We want to define order-raising operators $\hat{U}_i: \mathbb{R}P_i \rightarrow \mathbb{R}P_{i+1}$ and order-lowering operators $\hat{D}_i: \mathbb{R}P_i \rightarrow \mathbb{R}P_{i-1}$. Let us first consider just \hat{U}_i . The idea is to identify the basis element $v_{\mathfrak{o}}$ of $\mathbb{R}B_n^G$ with the basis element \mathfrak{o} of $\mathbb{R}P$, and to let $\hat{U}_i: \mathbb{R}P_i \rightarrow \mathbb{R}P_{i+1}$ correspond to the usual order-raising operator $U_i: \mathbb{R}(B_n)_i \rightarrow \mathbb{R}(B_n)_{i+1}$. More precisely, suppose that the order-raising operator U_i for B_n given by (4.2) satisfies

$$U_i(v_{\mathfrak{o}}) = \sum_{\mathfrak{o}' \in (B_n)_{i+1}/G} c_{\mathfrak{o}, \mathfrak{o}'} v_{\mathfrak{o}'}, \quad (5.2)$$

where $\mathfrak{o} \in (B_n)_i/G$. (Note that by Lemma 5.7, $U_i(v_{\mathfrak{o}})$ does indeed have the form given by (5.2).) Then define the linear operator $\hat{U}_i: \mathbb{R}((B_n)_i/G) \rightarrow \mathbb{R}((B_n)_{i+1}/G)$ by

$$\hat{U}_i(\mathfrak{o}) = \sum_{\mathfrak{o}' \in (B_n)_{i+1}/G} c_{\mathfrak{o}, \mathfrak{o}'} \mathfrak{o}'.$$

NOTE. We can depict the “transport of U_i to \hat{U}_i ” by a *commutative diagram*:

$$\begin{array}{ccc} (\mathbb{R}B_n)_i^G & \xrightarrow{U_i} & (\mathbb{R}B_n)_{i+1}^G \\ \cong \downarrow & & \downarrow \cong \\ \mathbb{R}(B_n/G)_i & \xrightarrow{\hat{U}_i} & \mathbb{R}(B_n/G)_{i+1} \end{array}$$

The arrows pointing down are the linear transformations induced by $v_{\mathfrak{o}} \mapsto \mathfrak{o}$. The map obtained by applying the top arrow followed by the rightmost down arrow is the same as applying the leftmost down arrow followed by the bottom arrow.

We claim that \hat{U}_i is order-raising. We need to show that if $c_{\mathfrak{o}, \mathfrak{o}'} \neq 0$, then $\mathfrak{o}' > \mathfrak{o}$ in B_n/G . Since $v_{\mathfrak{o}'} = \sum_{x' \in \mathfrak{o}'} x'$, the only way $c_{\mathfrak{o}, \mathfrak{o}'} \neq 0$ in (5.2) is for some $x' \in \mathfrak{o}'$ to satisfy $x' > x$ for some $x \in \mathfrak{o}$. But this is just what it means for $\mathfrak{o}' > \mathfrak{o}$, so \hat{U}_i is order-raising.

Now comes the heart of the argument. We want to show that \hat{U}_i is one-to-one for $i < n/2$. Now by Theorem 4.7, U_i is one-to-one for $i < n/2$. Thus the restriction of U_i to the subspace $\mathbb{R}(B_n)_i^G$ is one-to-one. (The restriction of a one-to-one function is always one-to-one.) But U_i and \hat{U}_i are exactly the same transformation, except for the names of the basis elements on which they act. Thus \hat{U}_i is also one-to-one for $i < n/2$.

An exactly analogous argument can be applied to D_i instead of U_i . We obtain one-to-one order-lowering operators $\hat{D}_i: \mathbb{R}(B_n)_i^G \rightarrow \mathbb{R}(B_n)_{i-1}^G$ for $i > n/2$. It follows from Proposition 4.4, Lemma 4.5, and equation (4.4) that B_n/G is rank-unimodal and Sperner, completing the proof. \square

We will consider two interesting applications of Theorem 5.8. For our first application, we let $n = \binom{m}{2}$ for some $m \geq 1$, and let $M = \{1, \dots, m\}$. Set $X = \binom{M}{2}$, the set of all two-element subsets of M . Think of the elements of X as (possible) edges of a simple graph with vertex set M . If B_X is the boolean algebra of all subsets of X (so B_X and B_n are isomorphic), then an element x of B_X is a collection of edges on the vertex set M , in other words, just a simple graph on M . Define a subgroup G of \mathfrak{S}_X as follows. Informally, G consists of all permutations of the edges $\binom{M}{2}$ that are induced from permutations of the vertices M . More precisely, if $\pi \in \mathfrak{S}_m$, then define $\hat{\pi} \in \mathfrak{S}_X$ by $\hat{\pi} \cdot \{i, j\} = \{\pi \cdot i, \pi \cdot j\}$. Thus G is isomorphic to \mathfrak{S}_m .

When are two graphs $x, y \in B_X$ in the same orbit of the action of G on B_X ? Since the elements of G just permute vertices, we see that x and y are in the same orbit if we can obtain x from y by permuting vertices. This is just what it means for two simple graphs x and y to be *isomorphic* — they are the same graph except for the names of the vertices (thinking of edges as pairs of vertices). Thus the elements of B_X/G are *isomorphism classes* of simple graphs on the vertex set M . In particular, $\#(B_X/G)$ is the number of nonisomorphic m -vertex simple graphs, and $\#(B_X/G)_i$ is the number of nonisomorphic such graphs with i edges. We have $x \leq y$ in B_X/G if there is some way of labelling the vertices of x and y so that every edge of x is an edge of y . Equivalently, some *spanning subgraph* of y (i.e., a subgraph of y with all the vertices of y) is isomorphic to x , as illustrated in Figure 5.1 for the case $m = 4$. Hence by Theorem 5.8 there follows the following result, which is by no means obvious and has no known non-algebraic proof.

5.9 Theorem. (a) Fix $m \geq 1$. Let p_i be the number of nonisomorphic simple graphs with m vertices and i edges. Then the sequence $p_0, p_1, \dots, p_{\binom{m}{2}}$




Figure 5.1: The poset B_X/G of nonisomorphic graphs with four vertices

is symmetric and unimodal.

(b) Let T be a collection of simple graphs with m vertices such that no element of T is isomorphic to a spanning subgraph of another element of T . Then $\#T$ is maximized by taking T to consist of all nonisomorphic simple graphs with $\lfloor \frac{1}{2} \binom{m}{2} \rfloor$ edges.

Our second example of the use of Theorem 5.8 is more subtle and will be the topic of the next chapter.

Digression: edge reconstruction. Much work has been done on “reconstruction problems,” that is, trying to reconstruct a mathematical structure such as a graph from some of its substructures. The most famous of such problems is *vertex reconstruction*: given a simple graph G on p vertices v_1, \dots, v_p , let G_i be the subgraph obtained by deleting vertex v_i (and all incident edges). Given the multiset $\{G_1, \dots, G_p\}$ of vertex-deleted subgraphs, can G be uniquely reconstructed? It is important to realize that the vertices are *unlabelled*, so given G_i we don’t know for any j which vertex is v_j . The famous *vertex reconstruction conjecture* (still open) states that for $p \geq 3$ any graph G can be reconstructed from the multiset $\{G_1, \dots, G_p\}$.

Here we will be concerned with *edge* reconstruction, another famous open problem. Given a simple graph G with edges e_1, \dots, e_q , let $H_i = G - e_i$, the graph obtained from G by removing the edge e_i .

Edge Reconstruction Conjecture. A simple graph G can be uniquely reconstructed from its number of vertices and the multiset $\{H_1, \dots, H_q\}$ of edge-deleted subgraphs.

NOTE. As in the case of vertex-reconstruction, the subgraphs H_i are unlabelled. The reason for including the number of vertices is that for any graph with no edges, we have $\{H_1, \dots, H_q\} = \emptyset$, so we need to specify the number of vertices to obtain G .

NOTE. It can be shown that if G can be vertex-reconstructed, then G can be edge reconstructed. Hence the vertex-reconstruction conjecture implies the edge-reconstruction conjecture.

The techniques developed above to analyze group actions on boolean algebra can be used to prove a special case of the edge-reconstruction conjecture. Note that a simple graph with p vertices has at most $\binom{p}{2}$ edges.

5.10 Theorem. *Let G be a simple graph with p vertices and $q > \frac{1}{2}\binom{p}{2}$ edges. Then G is edge-reconstructible.*

Proof. Let P_q be the set of all simple graphs with q edges on the vertex set $[p] = \{1, 2, \dots, p\}$, so $\#P_q = \binom{\binom{p}{2}}{q}$. Let $\mathbb{R}P_q$ denote the real vector space with basis P_q . Define a linear transformation $\psi_q: \mathbb{R}P_q \rightarrow \mathbb{R}P_{q-1}$ by

$$\psi_q(\Gamma) = \Gamma_1 + \dots + \Gamma_q,$$

where $\Gamma_1, \dots, \Gamma_q$ are the (labelled) graphs obtained from Γ by deleting a single edge. By Theorem 4.7, ψ_q is injective for $q > \frac{1}{2}\binom{p}{2}$. (Think of ψ_q as adding edges to the *complement* of Γ , i.e., the graph with vertex set $[p]$ and edge set $\binom{[p]}{2} - E(\Gamma)$.)

The symmetric group \mathfrak{S}_p acts on P_q by permuting the vertices, and hence acts on $\mathbb{R}P_q$, the real vector space with basis P_q . A basis for the fixed space $(\mathbb{R}P_q)^{\mathfrak{S}_p}$ consists of the distinct sums $\tilde{\Gamma} = \sum_{\pi \in \mathfrak{S}_p} \pi(\Gamma)$, where $\Gamma \in P_q$. We may identify $\tilde{\Gamma}$ with the *unlabelled* graph isomorphic to Γ , since $\tilde{\Gamma} = \tilde{\Gamma}'$ if and only if Γ and Γ' are isomorphic. Just as in the proof of Theorem 5.8, when we restrict ψ_q to $(\mathbb{R}P_q)^{\mathfrak{S}_p}$ for $q > \frac{1}{2}\binom{p}{2}$ we obtain an injection $\psi_q: (\mathbb{R}P_q)^{\mathfrak{S}_p} \rightarrow$

$(\mathbb{R}P_{q-1})^{\mathfrak{S}_p}$. In particular, for nonisomorphic unlabelled graphs $\tilde{\Gamma}, \tilde{\Gamma}'$ with p vertices, we have

$$\tilde{\Gamma}_1 + \cdots + \tilde{\Gamma}_q = \psi_q(\tilde{\Gamma}) \neq \psi_q(\tilde{\Gamma}') = \tilde{\Gamma}'_1 + \cdots + \tilde{\Gamma}'_q.$$

Hence the unlabelled graphs $\tilde{\Gamma}_1, \dots, \tilde{\Gamma}_q$ determine $\tilde{\Gamma}$, as desired. \square

Polynomials with real zeros. There are many techniques other than the linear algebra used to prove Theorem 5.8 for showing that sequences are unimodal. Here we will discuss a technique based on simple analysis (calculus) for showing that sequences are unimodal. In fact, we will consider some stronger properties than unimodality.

A sequence a_0, a_1, \dots, a_n of real numbers is called *logarithmically concave*, or *log-concave* for short, if $a_i^2 \geq a_{i-1}a_{i+1}$ for $1 \leq i \leq n-1$. We say that a_0, a_1, \dots, a_n is *strongly log-concave* if $b_i^2 \geq b_{i-1}b_{i+1}$ for $1 \leq i \leq n-1$, where $b_i = a_i / \binom{n}{i}$. Strong log-concavity is equivalent to [why?]

$$a_i^2 \geq \left(1 + \frac{1}{i}\right) \left(1 + \frac{1}{n-i}\right) a_{i-1}a_{i+1}, \quad 1 \leq i \leq n-1,$$

from which it follows that strong log-concavity implies log-concavity.

Assume now that each $a_i \geq 0$. Does log-concavity then imply unimodality? The answer is *no*, a counterexample being $1, 0, 0, 1$. However, only this type of counterexample can occur, as we now explain. We say that the sequence a_0, a_1, \dots, a_n has *no internal zeros* if whenever we have $i < j < k$, $a_i \neq 0$, and $a_k \neq 0$, then $a_j \neq 0$.

5.11 Proposition. *Let $\alpha = (a_0, a_1, \dots, a_n)$ be a sequence of nonnegative real numbers with no internal zeros. If α is log-concave, then α is unimodal.*

Proof. If there are only two values of j for which $a_j \neq 0$ then we always have $a_{i-1}a_{i+1} = 0$ so the conclusion is clear. Now assume that there are at least three values of j for which $a_j \neq 0$, and assume that the proposition is false. Then there exists $1 \leq i \leq n-1$ for which $a_{i-1} > a_i \leq a_{i+1}$ and $a_{i+1} > 0$, so $a_i^2 < a_{i-1}a_{i+1}$, a contradiction. \square

Now we come to a fundamental method for proving log-concavity.

5.12 Theorem (I. Newton). *Let*

$$P(x) = \sum_{i=0}^n b_i x^i = \sum_{i=0}^n \binom{n}{i} a_i x^i$$

be a real polynomial all of whose zeros are real numbers. Then the sequence b_0, b_1, \dots, b_n is strongly log-concave, or equivalently, the sequence a_0, a_1, \dots, a_n is log-concave. Moreover, if each $b_i \geq 0$ (so the zeros of $P(x)$ are nonpositive [why?]) then the sequence b_0, b_1, \dots, b_n has no internal zeros.

Proof. Let $\deg P(x) = m \leq n$. By the Fundamental Theorem of Algebra, $P(x)$ has exactly m real zeros, counting multiplicities. Suppose that α is a zero of multiplicity $r > 1$, so $P(x) = (x - \alpha)^r L(x)$ for some polynomial $L(x)$ satisfying $L(\alpha) \neq 0$. A simple computation shows that α is a zero of $P'(x)$ (the derivative of $P(x)$) of multiplicity $r - 1$. Moreover, if $\alpha < \beta$ are both zeros of $P(x)$, then Rolle's theorem shows that $P'(x)$ has a zero γ satisfying $\alpha < \gamma < \beta$. It follows [why?] that $P'(x)$ has at least $m - 1$ real zeros. Since $\deg P'(x) = m - 1$ we see that $P'(x)$ has exactly $m - 1$ real zeros and no other zeros.

Let $Q(x) = \frac{d^{i-1}}{dx^{i-1}}P(x)$. Thus $Q(x)$ is a polynomial of degree at most $m - i + 1$ with only real zeros. Let $R(x) = x^{m-i+1}Q(1/x)$, a polynomial of degree at most $m - i + 1$. The zeros of $R(x)$ are just reciprocals of those zeros of $Q(x)$ not equal to 0, with possible new zeros at 0. At any rate, all zeros of $R(x)$ are real. Now let $S(x) = \frac{d^{m-i-1}}{dx^{m-i-1}}R(x)$, a polynomial of degree at most two. By Rolle's theorem (with a suitable handling of multiple zeros as above), every zero of $S(x)$ is real. An explicit computation yields

$$S(x) = \frac{m!}{2}(a_{i-1}x^2 + 2a_i x + a_{i+1}).$$

If $a_{i-1} = 0$ then trivially $a_i^2 \geq a_{i-1}a_{i+1}$. Otherwise $S(x)$ is a quadratic polynomial. Since it has real zeros, its discriminant Δ is nonnegative. But

$$\Delta = (2a_i)^2 - 4a_{i-1}a_{i+1} = 4(a_i^2 - a_{i-1}a_{i+1}) \geq 0,$$

so the sequence a_0, a_1, \dots, a_n is log-concave as claimed.

It remains to show that if each $a_i \geq 0$ then the sequence a_0, a_1, \dots, a_n has no internal zeros. Suppose to the contrary that for some $i < j < k$ we have $a_i > 0, a_j = 0, a_k > 0$. By arguing as in the previous paragraph we will obtain a polynomial of the form $c + dx^{k-i}$ with only real zeros, where $c, d > 0$. But since $k - i \geq 2$ we have that every such polynomial has a nonreal zero [why?], a contradiction which completes the proof. \square

In order to give combinatorial applications of Theorem 5.12 we need to find polynomials with real zeros whose coefficients are of combinatorial interest. One such example appears in Exercise 9.8, based on the fact that the characteristic polynomial of a symmetric matrix has only real zeros.

Notes for Chapter 5

The techniques developed in this chapter had their origins in papers of L. H. Harper [54] and M. Pouzet and I. G. Rosenberg, [89]. The closest treatment to ours appears in a paper of R. Stanley [105]. This latter paper also contains the proof of Theorem 5.10 (edge reconstruction) given here. This result was first proved by L. Lovász [72] by an inclusion-exclusion argument. The condition $q > \frac{1}{2} \binom{p}{2}$ in Theorem 5.10 was improved to $q > p(\log_2 p - 1)$ by V. Müller [80] (generalizing the argument of Lovász) and by I. Krasikov and Y. Roditty [67] (generalizing the argument of Stanley).

For further information on Newton's Theorem 5.12, see e.g. G. H. Hardy, J. E. Littlewood, and G. Pólya [53, p. 52]. For a general survey on unimodality, log-concavity, etc., see Stanley [107], with a sequel by F. Brenti [13].


Chapter 6

Young diagrams and q -binomial coefficients

A *partition* λ of an integer $n \geq 0$ is a sequence $\lambda = (\lambda_1, \lambda_2, \dots)$ of integers $\lambda_i \geq 0$ satisfying $\lambda_1 \geq \lambda_2 \geq \dots$ and $\sum_{i \geq 1} \lambda_i = n$. Thus all but finitely many λ_i are equal to 0. Each $\lambda_i > 0$ is called a *part* of λ . We sometimes suppress 0's from the notation for λ , e.g., $(5, 2, 2, 1)$, $(5, 2, 2, 1, 0, 0, 0)$, and $(5, 2, 2, 1, 0, 0, \dots)$ all represent the same partition λ (of 10, with four parts). If λ is a partition of n , then we denote this by $\lambda \vdash n$ or $|\lambda| = n$.


6.1 Example. There are seven partitions of 5, namely (writing e.g. 221 as short for $(2, 2, 1)$): 5, 41, 32, 311, 221, 2111, and 11111.

The subject of partitions of integers has been extensively developed, but we will only be concerned here with a small part related to our previous discussion. Given positive integers m and n , let $L(m, n)$ denote the set of all partitions with at most m parts and with largest part at most n . For instance, $L(2, 3) = \{\emptyset, 1, 2, 3, 11, 21, 31, 22, 32, 33\}$. (Note that we are denoting by \emptyset the unique partition $(0, 0, \dots)$ with no parts.) If $\lambda = (\lambda_1, \lambda_2, \dots)$ and $\mu = (\mu_1, \mu_2, \dots)$ are partitions, then define $\lambda \leq \mu$ if $\lambda_i \leq \mu_i$ for all i . This makes the set of all partitions into a very interesting poset, denoted Y and called *Young's lattice* (named after the British mathematician Alfred Young, 1873–1940). (It is called “Young's lattice” rather than “Young's poset” because it turns out to have certain properties which define a *lattice*. However, these properties are irrelevant to us here, so we will not bother to define the notion of a lattice.) We will be looking at some properties of Y in Chapter 8. The partial ordering on Y , when restricted to $L(m, n)$, makes


Figure 6.1: The posets $L(1, 4)$, $L(2, 2)$, and $L(2, 3)$

$L(m, n)$ into a poset which also has some fascinating properties. Figure 6.1 shows $L(1, 4)$, $L(2, 2)$, and $L(2, 3)$, while Figure 6.2 shows $L(3, 3)$.


There is a nice geometric way of viewing partitions and the poset $L(m, n)$. The *Young diagram* (sometimes just called the *diagram*) of a partition λ is a left-justified array of squares, with λ_i squares in the i th row. For instance, the Young diagram of $(4, 3, 1, 1)$ looks like:




If dots are used instead of boxes, then the resulting diagram is called a *Ferrers diagram*. Thus the Ferrers diagram of $(4, 3, 1, 1)$ looks like



The advantage of Young diagrams over Ferrers diagrams is that we can


Figure 6.2: The poset $L(3, 3)$

put numbers in the boxes of a Young diagram, which we will do in Chapter 8. Observe that $L(m, n)$ is simply the set of Young diagrams D fitting in an $m \times n$ rectangle (where the upper-left (northwest) corner of D is the same as the northwest corner of the rectangle), ordered by inclusion. *We will always assume that when a Young diagram D is contained in a rectangle R , the northwest corners agree.* It is also clear from the Young diagram point of view that $L(m, n)$ and $L(n, m)$ are isomorphic partially ordered sets, the isomorphism being given by transposing the diagram (i.e., interchanging rows and columns). If λ has Young diagram D , then the partition whose diagram is D^t (the transpose of D) is called the *conjugate* of λ and is denoted λ' . For instance, $(4, 3, 1, 1)' = (4, 2, 2, 1)$, with diagram



6.2 Proposition. *The poset $L(m, n)$ is graded of rank mn and rank-symmetric. The rank of a partition λ is just $|\lambda|$ (the sum of the parts of λ or the number of squares in its Young diagram).*

Proof. As in the proof of Proposition 5.5, we leave to the reader everything except rank-symmetry. To show rank-symmetry, consider the complement $\bar{\lambda}$ of λ in an $m \times n$ rectangle R , i.e., all the squares of R except for λ . (Note that $\bar{\lambda}$ depends on m and n , and not just λ .) For instance, in $L(4, 5)$, the complement of $(4, 3, 1, 1)$ looks like




If we rotate the diagram of $\bar{\lambda}$ by 180° then we obtain the diagram of a partition $\tilde{\lambda} \in L(m, n)$ satisfying $|\lambda| + |\tilde{\lambda}| = mn$. This correspondence between λ and $\tilde{\lambda}$ shows that $L(m, n)$ is rank-symmetric. \square

Our main goal in this chapter is to show that $L(m, n)$ is rank-unimodal and Sperner. Let us write $p_i(m, n)$ as short for $p_i(L(m, n))$, the number of elements of $L(m, n)$ of rank i . Equivalently, $p_i(m, n)$ is the number of partitions of i with largest part at most n and with at most m parts, or, in other words, the number of distinct Young diagrams with i squares which fit inside an $m \times n$ rectangle (with the same northwest corner, as explained previously). Though not really necessary for our goal, it is nonetheless interesting to obtain some information on these numbers $p_i(m, n)$. First let us consider the total number $\#L(m, n)$ of elements in $L(m, n)$.

6.3 Proposition. *We have $\#L(m, n) = \binom{m+n}{m}$.*

Proof. We will give an elegant combinatorial proof, based on the fact that $\binom{m+n}{m}$ is equal to the number of sequences a_1, a_2, \dots, a_{m+n} , where each a_j is either N or E , and there are m N 's (and hence n E 's) in all. We will associate a Young diagram D contained in an $m \times n$ rectangle R with such a sequence as follows. Begin at the lower left-hand corner of R , and trace out the southeast boundary of D , ending at the upper right-hand corner of R . This is done by taking a sequence of unit steps (where each square of R is one unit in length), each step either north or east. Record the sequence of steps, using N for a step to the north and E for a step to the east.

Example. Let $m = 5$, $n = 6$, $\lambda = (4, 3, 1, 1)$. Then R and D are given by:



The corresponding sequence of N 's and E 's is $NENNEENENE$.

It is easy to see (left to the reader) that the above correspondence gives a bijection between Young diagrams D fitting in an $m \times n$ rectangle R , and sequences of m N 's and n E 's. Hence the number of diagrams is equal to $\binom{m+n}{m}$, the number of sequences. \square

We now consider how many elements of $L(m, n)$ have rank i . To this end, let q be an indeterminate; and given $j \geq 1$ define $(\mathbf{j}) = 1 + q + q^2 + \cdots + q^{j-1}$. Thus $(\mathbf{1}) = 1$, $(\mathbf{2}) = 1 + q$, $(\mathbf{3}) = 1 + q + q^2$, etc. Note that (\mathbf{j}) is a polynomial in q whose value at $q = 1$ is just j (denoted $(\mathbf{j})_{q=1} = j$). Next define $(\mathbf{j})! = (\mathbf{1})(\mathbf{2}) \cdots (\mathbf{j})$ for $j \geq 1$, and set $(\mathbf{0})! = 1$. Thus $(\mathbf{1})! = 1$, $(\mathbf{2})! = 1 + q$, $(\mathbf{3})! = (1 + q)(1 + q + q^2) = 1 + 2q + 2q^2 + q^3$, etc., and $(\mathbf{j})!_{q=1} = j!$. Finally define for $k \geq j \geq 0$,

$$\binom{\mathbf{k}}{\mathbf{j}} = \frac{(\mathbf{k})!}{(\mathbf{j})!(\mathbf{k} - \mathbf{j})!}.$$

The expression $\binom{\mathbf{k}}{\mathbf{j}}$ is called a q -binomial coefficient (or *Gaussian coefficient*). When q is regarded as a prime power rather than as an indeterminate, then Exercise 4.4 gives a definition of $\binom{n}{k}$ in terms of the field \mathbb{F}_q . In this chapter we have no need of this algebraic interpretation of $\binom{n}{k}$.

Since $(\mathbf{r})!_{q=1} = r!$, it is clear that

$$\binom{\mathbf{k}}{\mathbf{j}}_{q=1} = \binom{k}{j}.$$

One sometimes says that $\binom{\mathbf{k}}{\mathbf{j}}$ is a “ q -analogue” of the binomial coefficient $\binom{k}{j}$. There is no precise definition of a q -analogue $P(q)$ of some mathematical object P (such as a formula or definition). It should have the property that there is a reasonable way to interpret $P(1)$ as being P . Ideally $P(q)$ should have some interpretation involving \mathbb{F}_q when q is regarded as a prime power. The q -analogue of the set $\{1\}$ is the finite field \mathbb{F}_q , and the q -analogue of the set $[n] = \{1, 2, \dots, n\}$ is the vector space \mathbb{F}_q^n .

6.4 Example. We have $\binom{\mathbf{k}}{\mathbf{j}} = \binom{\mathbf{k}}{\mathbf{k} - \mathbf{j}}$ [why?]. Moreover,

$$\begin{aligned} \binom{\mathbf{k}}{\mathbf{0}} &= \binom{\mathbf{k}}{\mathbf{k}} = 1 \\ \binom{\mathbf{k}}{\mathbf{1}} &= \binom{\mathbf{k}}{\mathbf{k} - \mathbf{1}} = (\mathbf{k}) = 1 + q + q^2 + \cdots + q^{k-1} \\ \binom{\mathbf{4}}{\mathbf{2}} &= \frac{(\mathbf{4})(\mathbf{3})(\mathbf{2})(\mathbf{1})}{(\mathbf{2})(\mathbf{1})(\mathbf{2})(\mathbf{1})} = 1 + q + 2q^2 + q^3 + q^4 \\ \binom{\mathbf{5}}{\mathbf{2}} &= \binom{\mathbf{5}}{\mathbf{3}} = 1 + q + 2q^2 + 2q^3 + 2q^4 + q^5 + q^6. \end{aligned}$$

In the above example, $\binom{k}{j}$ was always a polynomial in q (and with non-negative integer coefficients). It is not obvious that this is always the case, but it will follow easily from the following lemma.

6.5 Lemma. *We have*

$$\binom{k}{j} = \binom{k-1}{j} + q^{k-j} \binom{k-1}{j-1}, \quad (6.1)$$

whenever $k \geq 1$, with the initial conditions $\binom{0}{0} = 1$, $\binom{k}{j} = 0$ if $j < 0$ or $j > k$ (the same initial conditions satisfied by the binomial coefficients $\binom{k}{j}$).

Proof. This is a straightforward computation. Specifically, we have

$$\begin{aligned} \binom{k-1}{j} + q^{k-j} \binom{k-1}{j-1} &= \frac{(k-1)!}{(j)!(k-1-j)!} + q^{k-j} \frac{(k-1)!}{(j-1)!(k-j)!} \\ &= \frac{(k-1)!}{(j-1)!(k-1-j)!} \left(\frac{1}{(j)} + \frac{q^{k-j}}{(k-j)} \right) \\ &= \frac{(k-1)!}{(j-1)!(k-1-j)!} \frac{(k-j) + q^{k-j}(j)}{(j)(k-j)} \\ &= \frac{(k-1)!}{(j-1)!(k-1-j)!} \frac{(k)}{(j)(k-j)} \\ &= \binom{k}{j}. \end{aligned}$$

□

Note that if we put $q = 1$ in (6.1) we obtain the well-known formula

$$\binom{k}{j} = \binom{k-1}{j} + \binom{k-1}{j-1},$$

which is just the recurrence defining Pascal's triangle. Thus equation (6.1) may be regarded as a q -analogue of the Pascal triangle recurrence.

We can regard equation (6.1) as a recurrence relation for the q -binomial coefficients. Given the initial conditions of Lemma 6.5, we can use (6.1) inductively to compute $\binom{k}{j}$ for any k and j . From this it is obvious by induction that the q -binomial coefficient $\binom{k}{j}$ is a polynomial in q with nonnegative integer coefficients. The following theorem gives an even stronger result, namely, an explicit combinatorial interpretation of the coefficients.

6.6 Theorem. Let $p_i(m, n)$ denote the number of elements of $L(m, n)$ of rank i . Then

$$\sum_{i \geq 0} p_i(m, n)q^i = \binom{m+n}{m}. \quad (6.2)$$

NOTE. The sum on the left-hand side is really a finite sum, since $p_i(m, n) = 0$ if $i > mn$.

Proof. Let $P(m, n)$ denote the left-hand side of (6.2). We will show that

$$P(0, 0) = 1, \text{ and } P(m, n) = 0 \text{ if } m < 0 \text{ or } n < 0 \quad (6.3)$$

$$P(m, n) = P(m, n-1) + q^n P(m-1, n). \quad (6.4)$$

Note that equations (6.3) and (6.4) completely determine $P(m, n)$. On the other hand, substituting $k = m+n$ and $j = m$ in (6.1) shows that $\binom{m+n}{m}$ also satisfies (6.4). Moreover, the initial conditions of Lemma 6.5 show that $\binom{m+n}{m}$ also satisfies (6.3). Hence (6.3) and (6.4) imply that $P(m, n) = \binom{m+n}{m}$, so to complete the proof we need only establish (6.3) and (6.4).

Equation (6.3) is clear, since $L(0, n)$ consists of a single point (the empty partition \emptyset), so $\sum_{i \geq 0} p_i(0, n)q^i = 1$; while $L(m, n)$ is empty (or undefined, if you prefer) if $m < 0$ or $n < 0$.

The crux of the proof is to show (6.4). Taking the coefficient of q^i of both sides of (6.4), we see [why?] that (6.4) is equivalent to

$$p_i(m, n) = p_i(m, n-1) + p_{i-n}(m-1, n). \quad (6.5)$$

Consider a partition $\lambda \vdash i$ whose Young diagram D fits in an $m \times n$ rectangle R . If D does not contain the upper right-hand corner of R , then D fits in an $m \times (n-1)$ rectangle, so there are $p_i(m, n-1)$ such partitions λ . If on the other hand D does contain the upper right-hand corner of R , then D contains the whole first row of R . When we remove the first row of R , we have left a Young diagram of size $i-n$ which fits in an $(m-1) \times n$ rectangle. Hence there are $p_{i-n}(m-1, n)$ such λ , and the proof follows [why?]. \square

Note that if we set $q = 1$ in (6.2), then the left-hand side becomes $\#L(m, n)$ and the right-hand side $\binom{m+n}{m}$, agreeing with Proposition 6.3.

As the reader may have guessed by now, the poset $L(m, n)$ is isomorphic to a quotient poset B_s/G for a suitable integer $s > 0$ and finite group G

acting on B_s . Actually, it is clear that we must have $s = mn$ since $L(m, n)$ has rank mn and in general B_s/G has rank s . What is not so clear is the right choice of G . To this end, let $R = R_{mn}$ denote an $m \times n$ rectangle of squares. For instance, R_{35} is given by the 15 squares of the diagram

We now define the group $G = G_{mn}$ as follows. It is a subgroup of the group \mathfrak{S}_R of all permutations of the squares of R . A permutation π in G is allowed to permute the elements in each row of R in any way, and then to permute the rows among themselves in any way. The elements of each row can be permuted in $n!$ ways, so since there are m rows there are a total of $n!^m$ permutations preserving the rows. Then the m rows can be permuted in $m!$ ways, so it follows that the order of G_{mn} is given by $m!n!^m$. The group G_{mn} is called the *wreath product* of \mathfrak{S}_n and \mathfrak{S}_m , denoted $\mathfrak{S}_n \wr \mathfrak{S}_m$ or $\mathfrak{S}_n \wr \mathfrak{S}_m$. However, we will not discuss the general theory of wreath products here.

6.7 Example. Suppose $m = 4$ and $n = 5$, with the boxes of R labelled as follows.

1	2	3	4	5
6	7	8	9	10
11	12	13	14	15
16	17	18	19	20

Then a typical permutation π in $G(4, 5)$ looks like

16	20	17	19	18
4	1	5	2	3
12	13	15	14	11
7	9	6	10	8

i.e., $\pi(16) = 1$, $\pi(20) = 2$, etc.

We have just defined a group G_{mn} of permutations of the set $R = R_{mn}$ of squares of an $m \times n$ rectangle. Hence G_{mn} acts on the boolean algebra B_R of all subsets of the set R . The next lemma describes the orbits of this action.

6.8 Lemma. *Every orbit \mathfrak{o} of the action of G_{mn} on B_R contains exactly one Young diagram D , i.e., exactly one subset $D \subseteq R$ such that D is left-justified, and if λ_i is the number of elements of D in row i of R , then $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_m$.*

Proof. Let S be a subset of R , and suppose that S has α_i elements in row i . If $\pi \in G_{mn}$ and $\pi \cdot S$ has β_i elements in row i , then β_1, \dots, β_m is just some permutation of $\alpha_1, \dots, \alpha_m$ [why?]. There is a unique ordering $\lambda_1, \dots, \lambda_m$ of $\alpha_1, \dots, \alpha_m$ satisfying $\lambda_1 \geq \cdots \geq \lambda_m$, so the only possible Young diagram D in the orbit $\pi \cdot S$ is the one of shape $\lambda = (\lambda_1, \dots, \lambda_m)$. It's easy to see that the Young diagram D_λ of shape λ is indeed in the orbit $\pi \cdot S$. For by permuting the elements in the rows of R we can left-justify the rows of S , and then by permuting the rows of R themselves we can arrange the row sizes of S to be in weakly decreasing order. Thus we obtain the Young diagram D_λ as claimed. \square

We are now ready for the main result of this chapter.

6.9 Theorem. *Set $R = R_{mn}$. Then the quotient poset B_R/G_{mn} is isomorphic to $L(m, n)$.*

Proof. Each element of B_R/G_{mn} contains a unique Young diagram D_λ by Lemma 6.8. Moreover, two different orbits cannot contain the same Young diagram D since orbits are disjoint. Thus the map $\varphi: B_R/G_{mn} \rightarrow L(m, n)$ defined by $\varphi(D_\lambda) = \lambda$ is a bijection (one-to-one and onto). We claim that in fact φ is an isomorphism of partially ordered sets. We need to show the following: let \mathfrak{o} and \mathfrak{o}^* be orbits of G_{mn} (i.e., elements of B_R/G_{mn}). Let D_λ and D_{λ^*} be the unique Young diagrams in \mathfrak{o} and \mathfrak{o}^* , respectively. Then there exist $D \in \mathfrak{o}$ and $D^* \in \mathfrak{o}^*$ satisfying $D \subseteq D^*$ if and only if $\lambda \leq \lambda^*$ in $L(m, n)$.

The “if” part of the previous sentence is clear, for if $\lambda \leq \lambda^*$ then $D_\lambda \subseteq D_{\lambda^*}$. So assume there exist $D \in \mathfrak{o}$ and $D^* \in \mathfrak{o}^*$ satisfying $D \subseteq D^*$. The lengths of the rows of D , written in decreasing order, are $\lambda_1, \dots, \lambda_m$, and similarly for D^* . Since each row of D is contained in a row of D^* , it follows that for each $1 \leq j \leq m$, D^* has at least j rows of size at least λ_j . Thus

the length λ_j^* of the j th largest row of D^* is at least as large as λ_j . In other words, $\lambda_j \leq \lambda_j^*$, as was to be proved. \square

Combining the previous theorem with Theorem 5.8 yields the following result.

6.10 Corollary. *The posets $L(m, n)$ are rank-symmetric, rank-unimodal, and Sperner.*

Note that the rank-symmetry and rank-unimodality of $L(m, n)$ can be rephrased as follows: the q -binomial coefficient $\binom{m+n}{m}$ has symmetric and unimodal coefficients. While rank-symmetry is easy to prove (see Proposition 6.2), the unimodality of the coefficients of $\binom{m+n}{m}$ is by no means apparent. It was first proved by J. Sylvester in 1878 by a proof similar to the one above, though stated in the language of the invariant theory of binary forms. For a long time it was an open problem to find a combinatorial proof that the coefficients of $\binom{m+n}{m}$ are unimodal. Such a proof would give an explicit injection (one-to-one function) $\mu: L(m, n)_i \rightarrow L(m, n)_{i+1}$ for $i < \frac{1}{2}mn$. (One difficulty in finding such maps μ is to make use of the hypothesis that $i < \frac{1}{2}mn$.) Finally around 1989 such a proof was found by K. M. O'Hara. However, O'Hara's proof has the defect that the maps μ are not order-matchings. Thus her proof does not prove that $L(m, n)$ is Sperner, but only that it's rank-unimodal. It is an outstanding open problem in algebraic combinatorics to find an explicit order-matching $\mu: L(m, n)_i \rightarrow L(m, n)_{i+1}$ for $i < \frac{1}{2}mn$.

Note that the Sperner property of $L(m, n)$ (together with the fact that the largest level is in the middle) can be stated in the following simple terms: the largest possible collection \mathcal{C} of Young diagrams fitting in an $m \times n$ rectangle such that no diagram in \mathcal{C} is contained in another diagram in \mathcal{C} is obtained by taking all the diagrams of size $\lfloor \frac{1}{2}mn \rfloor$. Although the statement of this fact requires almost no mathematics to understand, there is no known proof that doesn't use algebraic machinery. The several known algebraic proofs are all closely related, and the one we have given is the simplest. Corollary 6.10 is a good example of the efficacy of algebraic combinatorics.

An application to number theory. There is an interesting application of Corollary 6.10 to a number-theoretic problem. Fix a positive integer k . For a finite subset S of $\mathbb{R}^+ = \{\beta \in \mathbb{R} : \beta > 0\}$, and for a real number $\alpha > 0$,

define

$$f_k(S, \alpha) = \# \left\{ T \in \binom{S}{k} : \sum_{t \in T} t = \alpha \right\}.$$

In other words, $f_k(S, \alpha)$ is the number of k -element subsets of S whose elements sum to α . For instance, $f_3(\{1, 3, 4, 6, 7\}, 11) = 2$, since $1 + 3 + 7 = 1 + 4 + 6 = 11$.

Given positive integers $k < n$, our object is to maximize $f_k(S, \alpha)$ subject to the condition that $\#S = n$. We are free to choose both S and α , but k and n are fixed. Call this maximum value $h_k(n)$. Thus

$$h_k(n) = \max_{\substack{\alpha \in \mathbb{R}^+ \\ S \subseteq \mathbb{R}^+ \\ \#S = n}} f_k(S, \alpha).$$

What sort of behavior can we expect of the maximizing set S ? If the elements of S are “spread out,” say $S = \{1, 2, 4, 8, \dots, 2^{n-1}\}$, then all the subset sums of S are distinct. Hence for any $\alpha \in \mathbb{R}^+$ we have $f_k(S, \alpha) = 0$ or 1 . Similarly, if the elements of S are “unrelated” (e.g., linearly independent over the rationals, such as $S = \{1, \sqrt{2}, \sqrt{3}, \pi, \pi^2\}$), then again all subset sums are distinct and $f_k(S, \alpha) = 0$ or 1 . These considerations make it plausible that we should take $S = [n] = \{1, 2, \dots, n\}$ and then choose α appropriately. In other words, we are led to the conjecture that for any $S \in \binom{\mathbb{R}^+}{n}$ and $\alpha \in \mathbb{R}^+$, we have

$$f_k(S, \alpha) \leq f_k([n], \beta), \quad (6.6)$$

for some $\beta \in \mathbb{R}^+$ to be determined.

First let us evaluate $f_k([n], \alpha)$ for any α . This will enable us to determine the value of β in equation (6.6). Let $S = \{i_1, \dots, i_k\} \subseteq [n]$ with

$$1 \leq i_1 < i_2 < \dots < i_k \leq n, \quad i_1 + \dots + i_k = \alpha. \quad (6.7)$$

Let $j_r = i_r - r$. Then (since $1 + 2 + \dots + k = \binom{k+1}{2}$)

$$n - k \geq j_k \geq j_{k-1} \geq \dots \geq j_1 \geq 0, \quad j_1 + \dots + j_k = \alpha - \binom{k+1}{2}. \quad (6.8)$$

Conversely, given j_1, \dots, j_k satisfying (6.8) we can recover i_1, \dots, i_k satisfying (6.7). Hence $f_k([n], \alpha)$ is equal to the number of sequences j_1, \dots, j_k satisfying (6.8). Now let

$$\lambda(S) = (j_k, j_{k-1}, \dots, j_1).$$

Note that $\lambda(S)$ is a partition of the integer $\alpha - \binom{k+1}{2}$ with at most k parts and with largest part at most $n - k$. Thus

$$f_k([n], \alpha) = p_{\alpha - \binom{k+1}{2}}(k, n - k), \quad (6.9)$$

or equivalently,

$$\sum_{\alpha \geq \binom{k+1}{2}} f_k([n], \alpha) q^{\alpha - \binom{k+1}{2}} = \binom{n}{k}.$$

By the rank-unimodality (and rank-symmetry) of $L(n-k, k)$ (Corollary 6.10), the largest coefficient of $\binom{n}{k}$ is the middle one, that is, the coefficient of $\lfloor k(n-k)/2 \rfloor$. It follows that for fixed k and n , $f_k([n], \alpha)$ is maximized for $\alpha = \lfloor k(n-k)/2 \rfloor + \binom{k+1}{2} = \lfloor k(n+1)/2 \rfloor$. Hence the following result is plausible.

6.11 Theorem. *Let $S \in \binom{\mathbb{R}^+}{n}$, $\alpha \in \mathbb{R}^+$, and $k \in \mathbb{P}$. Then*

$$f_k(S, \alpha) \leq f_k([n], \lfloor k(n+1)/2 \rfloor).$$

Proof. Let $S = \{a_1, \dots, a_n\}$ with $0 < a_1 < \dots < a_n$. Let T and U be distinct k -element subsets of S with the same element sums, say $T = \{a_{i_1}, \dots, a_{i_k}\}$ and $U = \{a_{j_1}, \dots, a_{j_k}\}$ with $i_1 < i_2 < \dots < i_k$ and $j_1 < j_2 < \dots < j_k$. Define $T^* = \{i_1, \dots, i_k\}$ and $U^* = \{j_1, \dots, j_k\}$, so $T^*, U^* \in \binom{[n]}{k}$. The crucial observation is the following:

Claim. The elements $\lambda(T^*)$ and $\lambda(U^*)$ are incomparable in $L(k, n-k)$, i.e., neither $\lambda(T^*) \leq \lambda(U^*)$ nor $\lambda(U^*) \leq \lambda(T^*)$.

Proof of claim. Suppose not, say $\lambda(T^*) \leq \lambda(U^*)$ to be definite. Thus by definition of $L(k, n-k)$ we have $i_r - r \leq j_r - r$ for $1 \leq r \leq k$. Hence $i_r \leq j_r$ for $1 \leq r \leq k$, so also $a_{i_r} \leq a_{j_r}$ (since $a_1 < \dots < a_n$). But $a_{i_1} + \dots + a_{i_k} = a_{j_1} + \dots + a_{j_k}$ by assumption, so $a_{i_r} = a_{j_r}$ for all r . This contradicts the assumption that T and U are distinct and proves the claim.

It is now easy to complete the proof of Theorem 6.11. Suppose that S_1, \dots, S_r are distinct k -element subsets of S with the same element sums. By the claim, $\{\lambda(S_1^*), \dots, \lambda(S_r^*)\}$ is an antichain in $L(k, n-k)$. Hence r cannot exceed the size of the largest antichain in $L(k, n-k)$. By Theorem 6.6 and Corollary 6.10, the size of the largest antichain in $L(k, n-k)$ is given by

$p_{\lfloor k(n-k)/2 \rfloor}(k, n-k)$. By equation (6.9) this number is equal to $f_k([n], \lfloor k(n+1)/2 \rfloor)$. In other words,

$$r \leq f_k([n], \lfloor k(n+1)/2 \rfloor),$$

which is what we wanted to prove. \square

Note that an equivalent statement of Theorem 6.11 is that $h_k(n)$ is equal to the coefficient of $q^{\lfloor k(n-k)/2 \rfloor}$ in $\binom{n}{k}$ [why?].

Variation on a theme. Suppose that in Theorem 6.11 we do not want to specify the cardinality of the subsets of S . In other words, for any $\alpha \in \mathbb{R}$ and any finite subset $S \subset \mathbb{R}^+$, define

$$f(S, \alpha) = \#\{T \subseteq S : \sum_{t \in T} t = \alpha\}.$$

How large can $f(S, \alpha)$ be if we require $\#S = n$? Call this maximum value $h(n)$. Thus

$$h(n) = \max_{\substack{\alpha \in \mathbb{R}^+ \\ S \subset \mathbb{R}^+ \\ \#S = n}} f(S, \alpha). \quad (6.10)$$

For instance, if $S = \{1, 2, 3\}$ then $f(S, 3) = 2$ (coming from the subsets $\{1, 2\}$ and $\{3\}$). This is easily seen to be best possible, i.e., $h(3) = 2$.

We will find $h(n)$ in a manner analogous to the proof of Theorem 6.11. The big difference is that the relevant poset $M(n)$ is *not* of the form B_n/G , so we will have to prove the injectivity of the order-raising operator U_i from scratch. Our proofs will be somewhat sketchy; it shouldn't be difficult for the reader who has come this far to fill in the details.

Let $M(n)$ be the set of all subsets of $[n]$, with the ordering $A \leq B$ if the elements of A are $a_1 > a_2 > \dots > a_j$ and the elements of B are $b_1 > b_2 > \dots > b_k$, where $j \leq k$ and $a_i \leq b_i$ for $1 \leq i \leq j$. (The empty set \emptyset is the bottom element of $M(n)$.) Figure 6.3 shows $M(1)$, $M(2)$, $M(3)$, and $M(4)$.

It is easy to see that $M(n)$ is graded of rank $\binom{n+1}{2}$. The rank of the subset $T = \{a_1, \dots, a_k\}$ is

$$\text{rank}(T) = a_1 + \dots + a_k. \quad (6.11)$$

It follows [why?] that the rank-generating function of $M(n)$ is given by

$$F(M(n), q) = \sum_{i=0}^{\binom{n+1}{2}} (\#M(n)_i) q^i = (1+q)(1+q^2) \cdots (1+q^n).$$

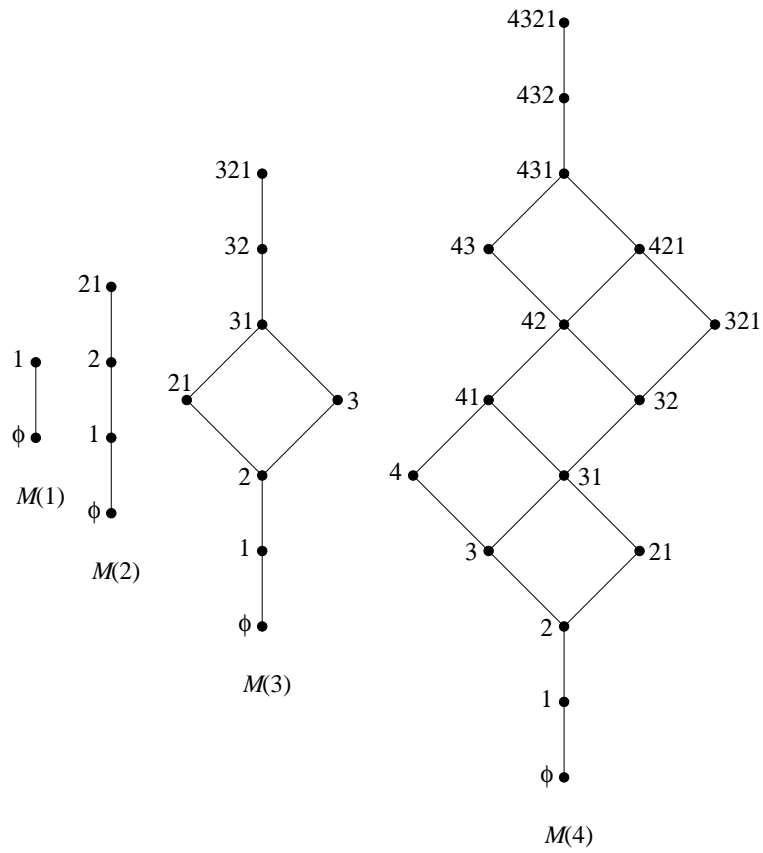


Figure 6.3: The posets $M(1)$, $M(2)$, $M(3)$ and $M(4)$

Define linear transformations

$$U_i: \mathbb{R}M(n)_i \rightarrow \mathbb{R}M(n)_{i+1}, \quad D_i: \mathbb{R}M(n)_i \rightarrow \mathbb{R}M(n)_{i-1}$$

by

$$U_i(x) = \sum_{\substack{y \in M(n)_{i+1} \\ x < y}} y, \quad x \in M(n)_i$$

$$D_i(x) = \sum_{\substack{v \in M(n)_{i-1} \\ v < x}} c(v, x)v, \quad x \in M(n)_i,$$

where the coefficient $c(v, x)$ is defined as follows. Let the elements of v be $a_1 > \cdots > a_j > 0$ and the elements of x be $b_1 > \cdots > b_k > 0$. Since x covers v , there is a unique r for which $a_r = b_r - 1$ (and $a_k = b_k$ for all other k). In the case $b_r = 1$ we set $a_r = 0$. (E.g., if x is given by $5 > 4 > 1$ and v by $5 > 4$, then $r = 3$ and $a_3 = 0$.) Set

$$c(v, x) = \begin{cases} \binom{n+1}{2}, & \text{if } a_r = 0 \\ (n - a_r)(n + a_r + 1), & \text{if } a_r > 0. \end{cases}$$

It is a straightforward computation (proof omitted) to obtain the commutation relation

$$D_{i+1}U_i - U_{i-1}D_i = \left(\binom{n+1}{2} - 2i \right) I_i, \quad (6.12)$$

where I_i denotes the identity linear transformation on $\mathbb{R}M(n)_i$. Clearly by definition U_i is order-raising. We want to show that U_i is injective (one-to-one) for $i < \frac{1}{2} \binom{n+1}{2}$. We can't argue as in the proof of Lemma 4.6 that $U_{i-1}D_i$ is semidefinite since the matrices of U_{i-1} and D_i are no longer transposes of one another. Instead we use the following result from linear algebra.

6.12 Lemma. *Let V and W be finite-dimensional vector spaces over a field. Let $A: V \rightarrow W$ and $B: W \rightarrow V$ be linear transformations. Then*

$$x^{\dim V} \det(AB - xI) = x^{\dim W} \det(BA - xI).$$

In other words, AB and BA have the same nonzero eigenvalues.

We can now prove the key linear algebraic result.

6.13 Lemma. *The linear transformation U_i is injective for $i < \frac{1}{2}\binom{n+1}{2}$ and surjective (onto) for $i \geq \frac{1}{2}\binom{n+1}{2}$.*

Proof. We prove by induction on i that $D_{i+1}U_i$ has positive real eigenvalues for $i < \frac{1}{2}\binom{n+1}{2}$. For $i = 0$ this is easy to check since $\dim \mathbb{R}M(n)_0 = 1$. Assume the induction hypothesis for some $i < \frac{1}{2}\binom{n+1}{2} - 1$, i.e., assume that D_iU_{i-1} has positive eigenvalues. By Lemma 6.12, $U_{i-1}D_i$ has nonnegative eigenvalues. By (6.12), we have

$$D_{i+1}U_i = U_{i-1}D_i + \left(\binom{n+1}{2} - 2i \right) I_i.$$

Thus the eigenvalues of $D_{i+1}U_i$ are $\binom{n+1}{2} - 2i$ more than those of $U_{i-1}D_i$. Since $\binom{n+1}{2} - 2i > 0$, it follows that $D_{i+1}U_i$ has positive eigenvalues. Hence it is invertible, so U_i is injective. Similarly (or by “symmetry”) U_i is surjective for $i \geq \frac{1}{2}\binom{n+1}{2}$. \square

The main result on the posets $M(n)$ now follows by a familiar argument.

6.14 Theorem. *The poset $M(n)$ is graded of rank $\binom{n+1}{2}$, rank-symmetric, rank-unimodal, and Sperner.*

Proof. We have already seen that $M(n)$ is graded of rank $\binom{n+1}{2}$ and rank-symmetric. By the previous lemma, U_i is injective for $i < \frac{1}{2}\binom{n+1}{2}$ and surjective for $i \geq \frac{1}{2}\binom{n+1}{2}$. The proof follows from Proposition 4.4 and Lemma 4.5. \square

NOTE. As a consequence of Theorem 6.14, the polynomial $F(M(n), q) = (1+q)(1+q^2)\cdots(1+q^n)$ has unimodal coefficients. No combinatorial proof of this fact is known, unlike the situation for $L(m, n)$ (where we mentioned the proof of O’Hara above).

We can now determine $h(n)$ (as defined by equation (6.10)) by an argument analogous to the proof of Theorem 6.11.

6.15 Theorem. *Let $S \in \binom{\mathbb{R}^+}{n}$ and $\alpha \in \mathbb{R}^+$. Then*

$$f(S, \alpha) \leq f\left([n], \left\lfloor \frac{1}{2}\binom{n+1}{2} \right\rfloor\right) = h(n).$$

Proof. Let $S = \{a_1, \dots, a_n\}$ with $0 < a_1 < \dots < a_n$. Let T and U be distinct subsets of S with the same element sums, say $T = \{a_{r_1}, \dots, a_{r_j}\}$ and $U = \{a_{s_1}, \dots, a_{s_k}\}$ with $r_1 < r_2 < \dots < r_j$ and $s_1 < s_2 < \dots < s_k$. Define $T^* = \{r_1, \dots, r_j\}$ and $U^* = \{s_1, \dots, s_k\}$, so $T^*, U^* \in M(n)$. The following fact is proved exactly in the same way as the analogous fact for $L(m, n)$ (the claim in the proof of Theorem 6.11) and will be omitted here.

Fact. The elements T^* and U^* are incomparable in $M(n)$, i.e., neither $T^* \leq U^*$ nor $U^* \leq T^*$.

It is now easy to complete the proof of Theorem 6.15. Suppose that S_1, \dots, S_t are distinct subsets of S with the same element sums. By the above fact, $\{S_1^*, \dots, S_t^*\}$ is an antichain in $M(n)$. Hence t cannot exceed the size of the largest antichain in $M(n)$. By Theorem 6.14, the size of the largest antichain in $M(n)$ is the size $p_{\lfloor \frac{1}{2} \binom{n+1}{2} \rfloor}$ of the middle rank. By equation (6.11) this number is equal to $f([n], \lfloor \frac{1}{2} \binom{n+1}{2} \rfloor)$. In other words,

$$t \leq f\left([n], \left\lfloor \frac{1}{2} \binom{n+1}{2} \right\rfloor\right),$$

which is what we wanted to prove. \square

NOTE. Theorem 6.15 is known as the *weak Erdős-Moser conjecture*. The original (strong) Erdős-Moser conjecture deals with the case $S \subset \mathbb{R}$ rather than $S \subset \mathbb{R}^+$. There is a difference between these two cases; for instance, $h(3) = 2$ (corresponding to $S = \{1, 2, 3\}$ and $\alpha = 3$), while the set $\{-1, 0, 1\}$ has *four* subsets whose elements sum to 0 (including the empty set). (Can you see where the proof of Theorem 6.15 breaks down if we allow $S \subset \mathbb{R}$?) The original Erdős-Moser conjecture asserts that if $\#S = 2m + 1$, then

$$f(S, \alpha) \leq f(\{-m, -m + 1, \dots, m\}, 0). \quad (6.13)$$

This result can be proved by a somewhat tricky modification of the proof given above for the weak case; see Exercise 6.5. No proof of the Erdős-Moser conjecture (weak or strong) is known other than the one indicated here (sometimes given in a more sophisticated context, as explained in the next Note).

NOTE. The key to the proof of Theorem 6.15 is the definition of U_i and D_i which gives the commutation relation (6.12). The reader may be wondering how anyone managed to discover these definitions (especially that of D_i).

In fact, the original proof of Theorem 6.15 was based on the representation theory of the orthogonal Lie algebra $\mathfrak{o}(2n+1, \mathbb{C})$. In this context, the definitions of U_i and D_i are built into the theory of the “principal subalgebras” of $\mathfrak{o}(2n+1, \mathbb{C})$. R. A. Proctor was the first to remove the representation theory from the proof and present it solely in terms of linear algebra.

Notes for Chapter 6

For an undergraduate level introduction to the theory of partitions, see Andrews and Eriksson [4]. A more extensive treatment is given by Andrews [3], while a brief introduction appears in [109, §1.8].

As already mentioned in the text, the rank-unimodality of $L(m, n)$, that is, of the coefficients of the q -binomial coefficient $\binom{m+n}{m}_q$, is due to J. J. Sylvester [114], with a combinatorial proof later given by K. M. O’Hara [83]. An explication of O’Hara’s work was given by D. Zeilberger [123].

The unimodality of the coefficients of the polynomial $(1+q)(1+q^2) \cdots (1+q^n)$ is implicit in the work of E. B. Dynkin [31][32, p. 332]. J. W. B. Hughes was the first to observe explicitly that this polynomial arises as a special case of Dynkin’s work. The Sperricity of $L(m, n)$ and $M(n)$, and a proof of the Erdős-Moser conjecture, were first given by Stanley [103]. It was mentioned in the text above that R. A. Proctor [90] was the first to remove the representation theory from the proof and present it solely in terms of linear algebra.


For two proofs of Lemma 6.12, see W. V. Parker [84] and J. Schmid [100].

Chapter 7

Enumeration under group action

In Chapters 5 and 6 we considered the quotient poset B_n/G , where G is a subgroup of the symmetric group \mathfrak{S}_n . If p_i is the number of elements of rank i of this poset, then the sequence p_0, p_1, \dots, p_n is rank-symmetric and rank-unimodal. Thus it is natural to ask whether there is some nice formula for the numbers p_i . For instance, in Theorem 5.9 p_i is the number of nonisomorphic graphs with m vertices (where $n = \binom{m}{2}$) and i edges; is there some nice formula for this number? For the group $G_{mn} = \mathfrak{S}_n \wr \mathfrak{S}_m$ of Theorem 6.6 we obtained a simple generating function for p_i (i.e., a formula for the rank-generating function $F(B_{mn}/G_{mn}, q) = \sum_i p_i q^i$), but this was a very special situation. In this chapter we will present a general theory for enumerating inequivalent objects subject to a group of symmetries, which will include a formula for the rank-generating functions $F(B_n/G, q)$. The chief architect of this theory is G. Pólya (though much of it was anticipated by J. H. Redfield) and hence is often called *Pólya's theory of enumeration* or just *Pólya theory*. See the references at the end of this chapter for further historical information.

Pólya theory is most easily understood in terms of “colorings” of some geometric or combinatorial object. For instance, consider a row of five squares:



In how many ways can we color the squares using n colors? Each square can

be colored any of the n colors, so there are n^5 ways in all. These colorings can be indicated as

A	B	C	D	E
---	---	---	---	---

where A, B, C, D, E are the five colors. Now assume that we are allowed to rotate the row of five squares 180° , and that two colorings are considered the same if one can be obtained from the other by such a rotation. (We may think that we have cut the row of five squares out of paper and colored them on one side.) We say that two colorings are *equivalent* if they are the same or can be transformed into one another by a 180° rotation. The first naive assumption is that every coloring is equivalent to exactly one other (besides itself), so the number of inequivalent colorings is $n^5/2$. Clearly this reasoning cannot be correct since $n^5/2$ is not always an integer! The problem, of course, is that some colorings stay the same when we rotate 180° . In fact, these are exactly the colorings

A	B	C	B	A
---	---	---	---	---

where A, B, C are any three colors. There are n^3 such colorings, so the total number of inequivalent colorings is given by

$$\begin{aligned} & \frac{1}{2}(\text{number of colorings which don't equal their } 180^\circ \text{ rotation}) \\ & + (\text{number of colorings which equal their } 180^\circ \text{ rotation}) \\ & = \frac{1}{2}(n^5 - n^3) + n^3 \\ & = \frac{1}{2}(n^5 + n^3). \end{aligned}$$

Pólya theory gives a systematic method for obtaining formulas of this sort for any underlying symmetry group.

The general setup is the following. Let X be a finite set, and G a subgroup of the symmetric group \mathfrak{S}_X . Think of G as a group of symmetries of X . Let C be another set (which may be infinite), which we think of as a set of “colors.” A *coloring* of X is a function $f: X \rightarrow C$. For instance, X could be the set of four squares of a 2×2 chessboard, labelled as follows:

1	2
3	4

Let $C = \{r, b, y\}$ (the colors red, blue, and yellow). A typical coloring of X would then look like

r	b
y	r

The above diagram thus indicates the function $f: X \rightarrow C$ given by $f(1) = r, f(2) = b, f(3) = y, f(4) = r$.

NOTE. We could work in the slightly greater generality of a group G acting on the set X , i.e., we are given a homomorphism $\varphi: G \rightarrow \mathfrak{S}_X$ that need not be injective. However, we then have a well-defined induced injective homomorphism $\psi: H \rightarrow \mathfrak{S}_X$, where $H = G/(\ker \varphi)$. The results obtained below for H are identical to those we get for G , so nothing is lost by assuming that φ is injective. In this case we can identify G with its image $\varphi(G)$.

We define two colorings f and g to be *equivalent* (or *G -equivalent*, when it is necessary to specify the group), denoted $f \sim g$ or $f \stackrel{G}{\sim} g$, if there exists an element $\pi \in G$ such that

$$g(\pi(x)) = f(x) \text{ for all } x \in X.$$

We may write this condition more succinctly as $g\pi = f$, where $g\pi$ denotes the composition of functions (from right to left). It is easy to check, using the fact that G is a group, that \sim is an equivalence relation. One should think that equivalent functions are the same “up to symmetry.”

7.1 Example. Let X be the 2×2 chessboard and $C = \{r, b, y\}$ as above. There are many possible choices of a symmetry group G , and this will affect when two colorings are equivalent. For instance, consider the following groups:

- G_1 consists of only the identity permutation $(1)(2)(3)(4)$.

- G_2 is the group generated by a vertical reflection. It consists of the two elements $(1)(2)(3)(4)$ (the identity element) and $(1, 2)(3, 4)$ (the vertical reflection).
- G_3 is the group generated by a reflection in the main diagonal. It consists of the two elements $(1)(2)(3)(4)$ (the identity element) and $(1)(4)(2, 3)$ (the diagonal reflection).
- G_4 is the group of all rotations of X . It is a cyclic group of order four with elements $(1)(2)(3)(4)$, $(1, 2, 4, 3)$, $(1, 4)(2, 3)$, and $(1, 3, 4, 2)$.
- G_5 is the dihedral group of all rotations and reflections of X . It has eight elements, namely, the four elements of G_4 and the four reflections $(1, 2)(3, 4)$, $(1, 3)(2, 4)$, $(1)(4)(2, 3)$, and $(2)(3)(1, 4)$.
- G_6 is the symmetric group of *all* 24 permutations of X . Although this is a perfectly valid group of symmetries, it no longer has any connection with the geometric representation of X as the squares of a 2×2 chessboard.

Consider the inequivalent colorings of X with two red squares, one blue square, and one yellow square, in each of the six cases above.

(G_1) There are twelve colorings in all with two red squares, one blue square, and one yellow square, and all are inequivalent under the trivial group (the group with one element). In general, whenever G is the trivial group then two colorings are equivalent if and only if they are the same [why?].

(G_2) There are now six inequivalent colorings, represented by

r	r	r	b	r	y	b	y	r	b	r	y
b	y	r	y	r	b	r	r	y	r	b	r

Each equivalence class contains two elements.

(G_3) Now there are seven classes, represented by

r	r	r	r	b	y	y	b	r	b	b	r	y	r
b	y	y	b	r	r	r	r	y	r	r	y	r	b

The first five classes contain two elements each and the last two classes only one element. Although G_2 and G_3 are isomorphic as abstract groups, as permutation groups they have a different structure. Specifically, the generator $(1, 2)(3, 4)$ of G_2 has two cycles of length two, while the generator $(1)(4)(2, 3)$ has two cycles of length one and one of length two. As we will see below, it is the lengths of the cycles of the elements of G that determine the sizes of the equivalence classes. This explains why the number of classes for G_2 and G_3 are different.

- (G_4) There are three classes, each with four elements. The size of each class is equal to the order of the group because none of the colorings have any symmetry with respect to the group, i.e., for any coloring f , the only group element π that fixes f (so $f\pi = f$) is the identity ($\pi = (1)(2)(3)(4)$).

<i>r</i>	<i>r</i>	<i>r</i>	<i>r</i>	<i>r</i>	<i>b</i>
<i>y</i>	<i>b</i>	<i>b</i>	<i>y</i>	<i>y</i>	<i>r</i>

- (G_5) Under the full dihedral group there are now two classes.

<i>r</i>	<i>r</i>	<i>r</i>	<i>b</i>
<i>b</i>	<i>y</i>	<i>y</i>	<i>r</i>

The first class has eight elements and the second four elements. In general, the size of a class is the index in G of the subgroup fixing some fixed coloring in that class [why?]. For instance, the subgroup fixing the second coloring above is $\{(1)(2)(3)(4), (1, 4)(2)(3)\}$, which has index four in the dihedral group of order eight.

- (G_6) Under the group \mathfrak{S}_4 of all permutations of the squares there is clearly only one class, with all twelve colorings. In general, for any set X if the group is the symmetric group \mathfrak{S}_X then two colorings are equivalent if and only if each color appears the same number of times [why?].

Our object in general is to count the number of equivalence classes of colorings which use each color a specified number of times. We will put the information into a *generating function* — a polynomial whose coefficients are the numbers we seek. Consider for example the set X , the group $G = G_5$ (the dihedral group), and the set $C = \{r, b, y\}$ of colors in Example 7.1 above. Let $\kappa(i, j, k)$ be the number of inequivalent colorings using red i times, blue j times, and yellow k times. Think of the colors r, b, y as *variables*, and form the polynomial

$$F_G(r, b, y) = \sum_{i+j+k=4} \kappa(i, j, k) r^i b^j y^k.$$

Note that we sum only over i, j, k satisfying $i + j + k = 4$ since a total of four colors will be used to color the four-element set X . The reader should check that

$$\begin{aligned} F_G(r, b, y) = & (r^4 + b^4 + y^4) + (r^3b + rb^3 + r^3y + ry^3 + b^3y + by^3) \\ & + 2(r^2b^2 + r^2y^2 + b^2y^2) + 2(r^2by + rb^2y + rby^2). \end{aligned}$$

For instance, the coefficient of r^2by is two because, as we have seen above, there are two inequivalent colorings using the colors r, r, b, y . Note that $F_G(r, b, y)$ is a *symmetric function* of the variables r, b, y (i.e., it stays the same if we permute the variables in any way), because insofar as counting inequivalent colorings goes, it makes no difference what *names* we give the colors. As a special case we may ask for the *total* number of inequivalent colorings with four colors. This is obtained by setting $r = b = y = 1$ in $F_G(r, b, y)$ [why?], yielding $F_G(1, 1, 1) = 3 + 6 + 2 \cdot 3 + 2 \cdot 3 = 21$.

What happens to the generating function F_G in the above example when we use the n colors r_1, r_2, \dots, r_n (which can be thought of as different shades of red)? Clearly all that matters are the *multiplicities* of the colors, without regard for their order. In other words, there are five cases: (a) all four colors the same, (b) one color used three times and another used once, (c) two colors used twice each, (d) one color used twice and two others once each, and (e) four colors used once each. These five cases correspond to the five partitions of 4, i.e., the five ways of writing 4 as a sum of positive integers without regard to order: $4, 3+1, 2+2, 2+1+1, 1+1+1+1$. Our generating

function becomes

$$\begin{aligned}
 F_G(r_1, r_2, \dots, r_n) &= \sum_i r_i^4 + \sum_{i \neq j} r_i^3 r_j \\
 &\quad + 2 \sum_{i < j} r_i^2 r_j^2 + 2 \sum_{\substack{i \neq j \\ i \neq k \\ j < k}} r_i^2 r_j r_k + 3 \sum_{i < j < k < l} r_i r_j r_k r_l,
 \end{aligned}$$

where the indices in each sum lie between 1 and n . If we set all variables equal to one (obtaining the total number of colorings with n colors), then simple combinatorial reasoning yields

$$\begin{aligned}
 F_G(1, 1, \dots, 1) &= n + n(n-1) + 2 \binom{n}{2} + 2n \binom{n-1}{2} + 3 \binom{n}{4} \\
 &= \frac{1}{8}(n^4 + 2n^3 + 3n^2 + 2n). \tag{7.1}
 \end{aligned}$$

Note that the polynomial (7.1) has the following description: the denominator 8 is the order of the group G_5 , and the coefficient of n^i in the numerator is just the number of permutations in G_5 with i cycles! For instance, the coefficient of n^2 is 3, and G_5 has the three elements $(1, 2)(3, 4)$, $(1, 3)(2, 4)$, and $(1, 4)(2, 3)$ with two cycles. We want to prove a general result of this nature.

The basic tool which we will use is a simple result from the theory of permutation groups known as *Burnside's lemma*. It was actually first proved by Cauchy when G is transitive (i.e., $|Y/G| = 1$ in Lemma 7.2 below) and by Frobenius in the general case, and is sometimes called the *Cauchy-Frobenius lemma*.

7.2 Lemma (Burnside's lemma). *Let Y be a finite set and G a subgroup of \mathfrak{S}_Y . For each $\pi \in G$, let*

$$\text{Fix}(\pi) = \{y \in Y : \pi(y) = y\},$$

so $\#\text{Fix}(\pi)$ is the number of cycles of length one in the permutation π . Let Y/G be the set of orbits of G . Then

$$|Y/G| = \frac{1}{\#G} \sum_{\pi \in G} \#\text{Fix}(\pi).$$

An equivalent form of Burnside's lemma is the statement that the average number of elements of Y fixed by an element of G is equal to the number of orbits. Before proceeding to the proof, let us consider an example.

7.3 Example. Let $Y = \{a, b, c, d\}$,

$$G = \{(a)(b)(c)(d), (a, b)(c, d), (a, c)(b, d), (a, d)(b, c)\},$$

and

$$G' = \{(a)(b)(c)(d), (a, b)(c)(d), (a)(b)(c, d), (a, b)(c, d)\}.$$

Both groups are isomorphic to $\mathbb{Z}_2 \times \mathbb{Z}_2$ (compare Example 5.1(c) and (d)). By Burnside's lemma the number of orbits of G is $\frac{1}{4}(4 + 0 + 0 + 0) = 1$. Indeed, given any two elements $i, j \in Y$, it is clear by inspection that there is a $\pi \in G$ (which happens to be unique) such that $\pi(i) = j$. On the other hand, the number of orbits of G' is $\frac{1}{4}(4 + 2 + 2 + 0) = 2$. Indeed, the two orbits are $\{a, b\}$ and $\{c, d\}$.

Proof of Burnside's lemma. For $y \in Y$ let $G_y = \{\pi \in G : \pi \cdot y = y\}$ (the set of permutations fixing y). Then

$$\begin{aligned} \frac{1}{\#G} \sum_{\pi \in G} \#\text{Fix}(\pi) &= \frac{1}{\#G} \sum_{\pi \in G} \sum_{\substack{y \in Y \\ \pi \cdot y = y}} 1 \\ &= \frac{1}{\#G} \sum_{y \in Y} \sum_{\substack{\pi \in G \\ \pi \cdot y = y}} 1 \\ &= \frac{1}{\#G} \sum_{y \in Y} \#G_y. \end{aligned}$$

Now (as in the proof of Lemma 5.6) the multiset of elements $\pi \cdot y$, $\pi \in G$, contains every element in the orbit Gy the same number of times, namely $\#G/\#Gy$ times. Thus y occurs $\#G/\#Gy$ times among the $\pi \cdot y$, so

$$\frac{\#G}{\#Gy} = \#G_y.$$

Therefore

$$\begin{aligned} \frac{1}{\#G} \sum_{\pi \in G} \#\text{Fix}(\pi) &= \frac{1}{\#G} \sum_{y \in Y} \frac{\#G}{\#Gy} \\ &= \sum_{y \in Y} \frac{1}{\#Gy}. \end{aligned}$$

How many times does a term $1/\#\mathcal{O}$ appear in the above sum, where \mathcal{O} is a fixed orbit? We are asking for the number of y such that $Gy = \mathcal{O}$. But $Gy = \mathcal{O}$ if and only if $y \in \mathcal{O}$, so $1/\#\mathcal{O}$ appears $\#\mathcal{O}$ times. Thus each orbit gets counted exactly once, so the above sum is equal to the number of orbits. \square

7.4 Example. How many inequivalent colorings of the vertices of a regular hexagon H are there using n colors, under cyclic symmetry? Let \mathcal{C}_n be the set of all n -colorings of H . Let G be the group of all permutations of \mathcal{C}_n which permute the colors cyclically, so $G \cong \mathbb{Z}_6$. We are asking for the number of orbits of G [why?]. We want to apply Burnside's lemma, so for each of the six elements σ of G we need to compute the number of colorings fixed by that element. Let π be a generator of G .

- $\sigma = 1$ (the identity): All n^6 colorings are fixed by σ .
- $\sigma = \pi, \pi^{-1}$: Only the n colorings with all colors equal are fixed.
- $\sigma = \pi^2, \pi^4$: Any coloring of the form $ababab$ is fixed (writing the colors linearly in the order they appear around the hexagon, starting at any fixed vertex). There are n choices for a and n for b , so n^2 colorings in all.
- $\sigma = \pi^3$: The fixed colorings are of the form $abcabc$, so n^3 in all.

Hence by Burnside's lemma, we have

$$\text{number of orbits} = \frac{1}{6}(n^6 + n^3 + 2n^2 + 2n).$$

The reader who has followed the preceding example will have no trouble understanding the following result.

7.5 Theorem. *Let G be a group of permutations of a finite set X . Then the number $N_G(n)$ of inequivalent (with respect to G) n -colorings of X is given by*

$$N_G(n) = \frac{1}{\#G} \sum_{\pi \in G} n^{c(\pi)}, \quad (7.2)$$

where $c(\pi)$ denotes the number of cycles of π .

Proof. Let π_n denote the action of $\pi \in G$ on the set \mathcal{C}_n of n -colorings of X . We want to determine the set $\text{Fix}(\pi_n)$, so that we can apply Burnside's lemma. Let C be the set of n colors. If $f: X \rightarrow C$ is a coloring fixed by π , then for all $x \in X$ we have

$$f(x) = \pi_n \cdot f(x) = f(\pi(x)).$$

Thus $f \in \text{Fix}(\pi_n)$ if and only if $f(x) = f(\pi(x))$. Hence $f(x) = f(\pi^k(x))$ for any $k \geq 1$ [why?]. The elements y of X of the form $\pi^k(x)$ for $k \geq 1$ are just the elements of the cycle of π containing x . Thus to obtain $f \in \text{Fix}(\pi_n)$, we should take the cycles $\sigma_1, \dots, \sigma_{c(\pi)}$ of π and color each element of σ_i the same color. There are n choices for each σ_i , so $n^{c(\pi)}$ colorings in all fixed by π . In other words, $\#\text{Fix}(\pi_n) = n^{c(\pi)}$, and the proof follows by Burnside's lemma. \square

We would now like not just to count the *total* number of inequivalent colorings with n colors, but more strongly to specify the number of occurrences of each color. We will need to use not just the number $c(\pi)$ of cycles of each $\pi \in G$, but rather the lengths of each of the cycles of π . Thus given a permutation π of an n -element set X , define the *type* of π to be

$$\text{type}(\pi) = (c_1, c_2, \dots, c_n),$$

where π has c_i i -cycles. For instance, if $\pi = (4, 7, 3, 8, 2, 10, 11, 1, 6, 9, 5)$, then

$$\begin{aligned} \text{type}(\pi) &= \text{type} (1, 4, 8)(2, 7, 11, 5)(3)(6, 10, 9) \\ &= (1, 0, 2, 1, 0, 0, 0, 0, 0, 0). \end{aligned}$$

Note that we always have $\sum_i i c_i = n$ [why?]. Define the *cycle indicator* of π to be the monomial

$$Z_\pi = z_1^{c_1} z_2^{c_2} \dots z_n^{c_n}.$$

(Many other notations are used for the cycle indicator. The use of Z_π comes from the German word *Zyklus* for cycle. The original paper of Pólya was written in German.) Thus for the example above, we have $Z_\pi = z_1 z_3^2 z_4$.

Now given a subgroup G of \mathfrak{S}_X , the *cycle indicator* (or *cycle index polynomial*) of G is defined by

$$Z_G = Z_G(z_1, \dots, z_n) = \frac{1}{\#G} \sum_{\pi \in G} Z_\pi.$$

Thus Z_G (also denoted P_G , $\text{Cyc}(G)$, etc.) is a polynomial in the variables z_1, \dots, z_n .

7.6 Example. If X consists of the vertices of a square and G is the group of rotations of X (a cyclic group of order 4), then

$$Z_G = \frac{1}{4}(z_1^4 + z_2^2 + 2z_4).$$

If reflections are also allowed (so G is the dihedral group of order 8), then

$$Z_G = \frac{1}{8}(z_1^4 + 3z_2^2 + 2z_1^2z_2 + 2z_4).$$

We are now ready to state the main result of this chapter.

7.7 Theorem (Pólya's theorem, 1937). *Let G be a group of permutations of the n -element set X . Let $C = \{r_1, r_2, \dots\}$ be a set of colors. Let $\kappa(i_1, i_2, \dots)$ be the number of inequivalent (under the action of G) colorings $f: X \rightarrow C$ such that color r_j is used i_j times. Define*

$$F_G(r_1, r_2, \dots) = \sum_{i_1, i_2, \dots} \kappa(i_1, i_2, \dots) r_1^{i_1} r_2^{i_2} \dots$$

(Thus F_G is a polynomial or a power series in the variables r_1, r_2, \dots , depending on whether or not C is finite or infinite.) Then

$$F_G(r_1, r_2, \dots) =$$

$$Z_G(r_1 + r_2 + r_3 + \dots, r_1^2 + r_2^2 + r_3^2 + \dots, \dots, r_1^j + r_2^j + r_3^j + \dots, \dots).$$

(In other words, substitute $\sum_i r_i^j$ for z_j in Z_G .)

Before giving the proof let us consider an example.

7.8 Example. Suppose that in Example 7.6 our set of colors is $C = \{a, b, c, d\}$, and that we take G to be the group of cyclic symmetries. Then

$$\begin{aligned} F_G(a, b, c, d) &= \frac{1}{4}((a + b + c + d)^4 + (a^2 + b^2 + c^2 + d^2)^2 + 2(a^4 + b^4 + c^4 + d^4)) \\ &= (a^4 + \dots) + (a^3b + \dots) + 2(a^2b^2 + \dots) + 3(a^2bc + \dots) + 6abcd. \end{aligned}$$

An expression such as $(a^2b^2 + \dots)$ stands for the sum of all monomials in the variables a, b, c, d with exponents 2, 2, 0, 0 (in some order). The coefficient of all such monomials is 2, indicating two inequivalent colorings using one color

twice and another color twice. If instead G were the full dihedral group, we would get

$$\begin{aligned} F_G(a, b, c, d) &= \frac{1}{8} \left((a+b+c+d)^4 + 3(a^2+b^2+c^2+d^2)^2 \right. \\ &\quad \left. + 2(a+b+c+d)^2(a^2+b^2+c^2+d^2) + 2(a^4+b^4+c^4+d^4) \right) \\ &= (a^4 + \cdots) + (a^3b + \cdots) + 2(a^2b^2 + \cdots) + 2(a^2bc + \cdots) + 3abcd. \end{aligned}$$

Proof of Pólya's theorem. Let $\#X = t$ and $i_1 + i_2 + \cdots = t$, where each $i_j \geq 0$. Let $\mathbf{i} = (i_1, i_2, \dots)$, and let $\mathcal{C}_{\mathbf{i}}$ denote the set of all colorings of X with color r_j used i_j times. The group G acts on $\mathcal{C}_{\mathbf{i}}$, since if $f \in \mathcal{C}_{\mathbf{i}}$ and $\pi \in G$, then $\pi \cdot f \in \mathcal{C}_{\mathbf{i}}$. (“Rotating” a colored object does not change how many times each color appears.) Let $\pi_{\mathbf{i}}$ denote the action of π on $\mathcal{C}_{\mathbf{i}}$. We want to apply Burnside's lemma to compute the number of orbits, so we need to find $\#\text{Fix}(\pi_{\mathbf{i}})$.

In order for $f \in \text{Fix}(\pi_{\mathbf{i}})$, we must color X so that (a) in any cycle of π , all the elements get the same color, and (b) the color r_j appears i_j times. Consider the product

$$H_{\pi} = \prod_j (r_1^j + r_2^j + \cdots)^{c_j(\pi)},$$

where $c_j(\pi)$ is the number of j -cycles (cycles of length j) of π . When we expand this product as a sum of monomials $r_1^{j_1} r_2^{j_2} \cdots$, we get one of these monomials by choosing a term r_k^j from each factor of H_{π} and multiplying these terms together. Choosing r_k^j corresponds to coloring all the elements of some j -cycle with r_k . Since a factor $r_1^j + r_2^j + \cdots$ occurs precisely $c_j(\pi)$ times in H_{π} , choosing a term r_k^j from every factor corresponds to coloring X so that every cycle is monochromatic (i.e., all the elements of that cycle get the same color). The product of these terms r_k^j will be the monomial $r_1^{j_1} r_2^{j_2} \cdots$, where we have used color r_k a total of j_k times. It follows that the coefficient of $r_1^{i_1} r_2^{i_2} \cdots$ in H_{π} is equal to $\#\text{Fix}(\pi_{\mathbf{i}})$. Thus

$$H_{\pi} = \sum_{\mathbf{i}} \#\text{Fix}(\pi_{\mathbf{i}}) r_1^{i_1} r_2^{i_2} \cdots. \quad (7.3)$$

Now sum both sides of (7.3) over all $\pi \in G$ and divide by $\#G$. The left-hand side becomes

$$\frac{1}{\#G} \sum_{\pi \in G} \prod_j (r_1^j + r_2^j + \cdots)^{c_j(\pi)} = Z_G(r_1 + r_2 + \cdots, r_1^2 + r_2^2 + \cdots, \dots).$$

On the other hand, the right-hand side becomes

$$\sum_{\mathbf{i}} \left[\frac{1}{\#G} \sum_{\pi \in G} \#\text{Fix}(\pi_{\mathbf{i}}) \right] r_1^{i_1} r_2^{i_2} \cdots .$$

By Burnside's lemma, the expression in brackets is just the number of orbits of $\pi_{\mathbf{i}}$ acting on $\mathcal{C}_{\mathbf{i}}$, i.e., the number of inequivalent colorings using color r_j a total of i_j times, as was to be proved. \square

7.9 Example (necklaces). A *necklace* of length ℓ is a circular arrangement of ℓ (colored) beads. Two necklaces are considered the same if they are cyclic rotations of one another. Let X be a set of ℓ (uncolored) beads, say $X = \{1, 2, \dots, \ell\}$. Regarding the beads as being placed equidistantly on a circle in the order $1, 2, \dots, \ell$, let G be the cyclic group of rotations of X . Thus if π is the cycle $(1, 2, \dots, \ell)$, then $G = \{1, \pi, \pi^2, \dots, \pi^{\ell-1}\}$. For example, if $\ell = 6$ then the elements of G are

$$\begin{aligned} \pi^0 &= (1)(2)(3)(4)(5)(6) \\ \pi &= (1, 2, 3, 4, 5, 6) \\ \pi^2 &= (1, 3, 5)(2, 4, 6) \\ \pi^3 &= (1, 4)(2, 5)(3, 6) \\ \pi^4 &= (1, 5, 3)(2, 6, 4) \\ \pi^5 &= (1, 6, 5, 4, 3, 2). \end{aligned}$$

In general, if d is the greatest common divisor of m and ℓ (denoted $d = \gcd(m, \ell)$), then π^m has d cycles of length ℓ/d . An integer m satisfies $1 \leq m \leq \ell$ and $\gcd(m, \ell) = d$ if and only if $1 \leq m/d \leq \ell/d$ and $\gcd(m/d, \ell/d) = 1$. Hence the number of such integers m is given by the Euler phi-function (or totient function) $\phi(\ell/d)$, which by definition is equal to the number of integers $1 \leq i \leq \ell/d$ such that $\gcd(i, \ell/d) = 1$. As an aside, recall that $\phi(k)$ can be computed by the formula

$$\phi(k) = k \prod_{\substack{p|k \\ p \text{ prime}}} \left(1 - \frac{1}{p}\right). \quad (7.4)$$

For instance, $\phi(1000) = 1000(1 - \frac{1}{2})(1 - \frac{1}{5}) = 400$. Putting all this together gives the following formula for the cycle enumerator $Z_G(z_1, \dots, z_{\ell})$:

$$Z_G(z_1, \dots, z_{\ell}) = \frac{1}{\ell} \sum_{d|\ell} \phi(\ell/d) z_{\ell/d}^d,$$

or (substituting ℓ/d for d),

$$Z_G(z_1, \dots, z_\ell) = \frac{1}{\ell} \sum_{d|\ell} \phi(d) z_d^{\ell/d}.$$

There follows from Pólya's theorem the following result (originally proved by P. A. MacMahon (1854–1929) before Pólya discovered his general result).

7.10 Theorem. (a) *The number $N_\ell(n)$ of n -colored necklaces of length ℓ is given by*

$$N_\ell(n) = \frac{1}{\ell} \sum_{d|\ell} \phi(\ell/d) n^d. \quad (7.5)$$

(b) *We have*

$$F_G(r_1, r_2, \dots) = \frac{1}{\ell} \sum_{d|\ell} \phi(d) (r_1^d + r_2^d + \dots)^{\ell/d}.$$

NOTE: (b) reduces to (a) if $r_1 = r_2 = \dots = 1$. Moreover, since clearly $N_\ell(1) = 1$, putting $n = 1$ in (7.5) yields the well-known identity

$$\sum_{d|\ell} \phi(\ell/d) = \ell.$$

What if we are allowed to flip necklaces over, not just rotate them? Now the group becomes the dihedral group of order 2ℓ , and the corresponding inequivalent colorings are called *dihedral necklaces*. We leave to the reader to work out the cycle enumerators

$$\frac{1}{2\ell} \left(\sum_{d|\ell} \phi(d) z_d^{\ell/d} + m z_1^2 z_2^{m-1} + m z_2^m \right), \quad \text{if } \ell = 2m$$

$$\frac{1}{2\ell} \left(\sum_{d|\ell} \phi(d) z_d^{\ell/d} + \ell z_1 z_2^m \right), \quad \text{if } \ell = 2m + 1.$$

7.11 Example. Let $G = \mathfrak{S}_\ell$, the group of all permutations of $\{1, 2, \dots, \ell\} = X$. Thus for instance

$$Z_{\mathfrak{S}_3}(z_1, z_2, z_3) = \frac{1}{6} (z_1^3 + 3z_1 z_2 + 2z_3)$$

$$Z_{\mathfrak{S}_4}(z_1, z_2, z_3, z_4) = \frac{1}{24} (z_1^4 + 6z_1^2 z_2 + 3z_2^2 + 8z_1 z_3 + 6z_4).$$

It is easy to count the number of inequivalent colorings in \mathcal{C}_i . If two colorings of X use each color the same number of times, then clearly there is *some* permutation of X which sends one of the colorings to the other. Hence \mathcal{C}_i consists of a single orbit. Thus

$$F_{\mathfrak{S}_\ell}(r_1, r_2, \dots) = \sum_{i_1+i_2+\dots=\ell} r_1^{i_1} r_2^{i_2} \dots,$$

the sum of all monomials of degree ℓ .

To count the total number of inequivalent n -colorings, note that

$$\sum_{\ell \geq 0} F_{\mathfrak{S}_\ell}(r_1, r_2, \dots) x^\ell = \frac{1}{(1-r_1x)(1-r_2x)\dots}, \quad (7.6)$$

since if we expand each factor on the right-hand side into the series $\sum_{j \geq 0} r_i^j x^j$ and multiply, the coefficient of x^ℓ will just be the sum of all monomials of degree ℓ . For fixed n , let $f_n(\ell)$ denote the number of inequivalent n -colorings of X . Since $f_n(\ell) = F_{\mathfrak{S}_\ell}(1, 1, \dots, 1)$ (n 1's in all), there follows from (7.6) that

$$\sum_{\ell \geq 0} f_n(\ell) x^\ell = \frac{1}{(1-x)^n}.$$

The right-hand side can be expanded (e.g. by Taylor's theorem) as

$$\frac{1}{(1-x)^n} = \sum_{\ell \geq 0} \binom{n+\ell-1}{\ell} x^\ell.$$

Hence

$$f_n(\ell) = \binom{n+\ell-1}{\ell}.$$

It is natural to ask whether there might be a more direct proof of such a simple result. This is actually a standard result in elementary enumerative combinatorics. For fixed ℓ and n we want the number of solutions to $i_1 + i_2 + \dots + i_n = \ell$ in nonnegative integers. Suppose that we arrange $n-1$ vertical bars and ℓ dots in a line. There are $\binom{n+\ell-1}{\ell}$ such arrangements since there are a total of $n+\ell-1$ positions, and we choose ℓ of them in which to place a dot. An example of such an arrangement for $\ell = 8$ and $n = 7$ is

$$| \quad | \quad \bullet \bullet \quad | \quad \bullet \quad | \quad | \quad \bullet \bullet \bullet \quad | \quad \bullet \bullet$$

The number of dots in each “compartment,” read from left to right, gives the numbers i_1, \dots, i_n . For the example above, we get $(i_1, \dots, i_7) = (0, 0, 2, 1, 0, 3, 2)$. Since this correspondence between solutions to $i_1 + i_2 + \dots + i_n = \ell$ and arrangements of bars and dots is clearly a bijection, we get $\binom{n+\ell-1}{\ell}$ solutions as claimed.

Recall (Theorem 7.5) that the number of inequivalent n -colorings of X (with respect to any group G of permutations of X) is given by

$$\frac{1}{\#G} \sum_{\pi \in G} n^{c(\pi)},$$

where $c(\pi)$ denotes the number of cycles of π . Hence for $G = \mathfrak{S}_\ell$ we get the identity

$$\begin{aligned} \frac{1}{\ell!} \sum_{\pi \in \mathfrak{S}_\ell} n^{c(\pi)} &= \binom{n+\ell-1}{\ell} \\ &= \frac{1}{\ell!} n(n+1)(n+2) \cdots (n+\ell-1). \end{aligned}$$

Multiplying by $\ell!$ yields

$$\sum_{\pi \in \mathfrak{S}_\ell} n^{c(\pi)} = n(n+1)(n+2) \cdots (n+\ell-1). \quad (7.7)$$

Equivalently [why?], if we define $c(\ell, k)$ to be the number of permutations in \mathfrak{S}_ℓ with k cycles (called a *signless Stirling number of the first kind*), then

$$\sum_{k=1}^{\ell} c(\ell, k) x^k = x(x+1)(x+2) \cdots (x+\ell-1).$$

For instance, $x(x+1)(x+2)(x+3) = x^4 + 6x^3 + 11x^2 + 6x$, so (taking the coefficient of x^2) eleven permutations in \mathfrak{S}_4 have two cycles, namely, $(123)(4)$, $(132)(4)$, $(124)(3)$, $(142)(3)$, $(134)(2)$, $(143)(2)$, $(234)(1)$, $(243)(1)$, $(12)(34)$, $(13)(24)$, $(14)(23)$.

Although it was easy to compute the generating function $F_{\mathfrak{S}_\ell}(r_1, r_2, \dots)$ directly without the necessity of computing the cycle indicator $Z_{\mathfrak{S}_\ell}(z_1, \dots, z_\ell)$, we can still ask whether there is a formula of some kind for this polynomial. First we determine explicitly its coefficients.

7.12 Theorem. Let $\sum ic_i = \ell$. The number of permutations $\pi \in \mathfrak{S}_\ell$ with c_i cycles of length i (or equivalently, the coefficient of $z_1^{c_1} z_2^{c_2} \cdots$ in $\ell! Z_{\mathfrak{S}_\ell}(z_1, \dots, z_\ell)$) is equal to $\ell! / 1^{c_1} c_1! 2^{c_2} c_2! \cdots$.

Example. The number of permutations in \mathfrak{S}_{15} with three 1-cycles, two 2-cycles, and two 4-cycles is $15! / 1^3 \cdot 3! \cdot 2^2 \cdot 2! \cdot 4^2 \cdot 2! = 851,350,500$.

Proof of Theorem 7.12. Fix $\mathbf{c} = (c_1, c_2, \dots)$ and let $X_{\mathbf{c}}$ be the set of all permutations $\pi \in \mathfrak{S}_\ell$ with c_i cycles of length i . Given a permutation $\sigma = a_1 a_2 \cdots a_\ell$ in \mathfrak{S}_ℓ , construct a permutation $f(\sigma) \in X_{\mathbf{c}}$ as follows. Let the 1-cycles of $f(\sigma)$ be $(a_1), (a_2), \dots, (a_{c_1})$. Then let the 2-cycles of $f(\sigma)$ be $(a_{c_1+1}, a_{c_1+2}), (a_{c_1+3}, a_{c_1+4}), \dots, (a_{c_1+2c_2-1}, a_{c_1+2c_2})$. Then let the 3-cycles of $f(\sigma)$ be $(a_{c_1+2c_2+1}, a_{c_1+2c_2+2}, a_{c_1+2c_2+3}), (a_{c_1+2c_2+4}, a_{c_1+2c_2+5}, a_{c_1+2c_2+6}), \dots, (a_{c_1+2c_2+3c_3-2}, a_{c_1+2c_2+3c_3-1}, a_{c_1+2c_2+3c_3})$, etc., continuing until we reach a_ℓ and have produced a permutation in $X_{\mathbf{c}}$. For instance, if $\ell = 11, c_1 = 3, c_2 = 2, c_4 = 1$, and $\sigma = 4, 9, 6, 11, 7, 1, 3, 8, 10, 2, 5$, then

$$f(\sigma) = (4)(9)(6)(11, 7)(1, 3)(8, 10, 2, 5).$$

We have defined a function $f: \mathfrak{S}_\ell \rightarrow X_{\mathbf{c}}$. Given $\pi \in X_{\mathbf{c}}$, what is $\#f^{-1}(\pi)$, the number of permutations sent to π by f ? A cycle of length i can be written in i ways, namely,

$$(b_1, b_2, \dots, b_i) = (b_2, b_3, \dots, b_i, b_1) = \cdots = (b_i, b_1, b_2, \dots, b_{i-1}).$$

Moreover, there are $c_i!$ ways to order the c_i cycles of length i . Hence

$$\#f^{-1}(\pi) = c_1! c_2! c_3! \cdots 1^{c_1} 2^{c_2} 3^{c_3} \cdots,$$

the same number for any $\pi \in X_{\mathbf{c}}$. It follows that

$$\begin{aligned} \#X_{\mathbf{c}} &= \frac{\#\mathfrak{S}_\ell}{c_1! c_2! \cdots 1^{c_1} 2^{c_2} \cdots} \\ &= \frac{\ell!}{c_1! c_2! \cdots 1^{c_1} 2^{c_2} \cdots}, \end{aligned}$$

as was to be proved. \square

As for the polynomial $Z_{\mathfrak{S}_\ell}$ itself, we have the following result. Write $\exp y = e^y$.

7.13 Theorem. *We have*

$$\sum_{\ell \geq 0} Z_{\mathfrak{S}_\ell}(z_1, z_2, \dots) x^\ell = \exp\left(z_1 x + z_2 \frac{x^2}{2} + z_3 \frac{x^3}{3} + \dots\right).$$

Proof. There are some sophisticated ways to prove this theorem which “explain” why the exponential function appears, but we will be content here with a “naive” computational proof. Write

$$\begin{aligned} & \exp\left(z_1 x + z_2 \frac{x^2}{2} + z_3 \frac{x^3}{3} + \dots\right) \\ &= e^{z_1 x} \cdot e^{z_2 \frac{x^2}{2}} \cdot e^{z_3 \frac{x^3}{3}} \dots \\ &= \left(\sum_{n \geq 0} \frac{z_1^n x^n}{n!}\right) \left(\sum_{n \geq 0} \frac{z_2^n x^{2n}}{2^n n!}\right) \left(\sum_{n \geq 0} \frac{z_3^n x^{3n}}{3^n n!}\right) \dots \end{aligned}$$

When we multiply this product out, the coefficient of $z_1^{c_1} z_2^{c_2} \dots x^\ell$, where $\ell = c_1 + 2c_2 + \dots$, is given by

$$\frac{1}{1^{c_1} c_1! 2^{c_2} c_2! \dots} = \frac{1}{\ell!} \left(\frac{\ell!}{1^{c_1} c_1! 2^{c_2} c_2! \dots}\right).$$

By Theorem 7.12 this is just the coefficient of $z_1^{c_1} z_2^{c_2} \dots$ in $Z_{\mathfrak{S}_\ell}(z_1, z_2, \dots)$, as was to be proved. \square

As a check of Theorem 7.13, set each $z_i = n$ to obtain

$$\begin{aligned} \sum_{\ell \geq 0} Z_{\mathfrak{S}_\ell}(n, n, \dots) x^\ell &= \exp\left(nx + n \frac{x^2}{2} + n \frac{x^3}{3} + \dots\right) \\ &= \exp\left(n\left(x + \frac{x^2}{2} + \frac{x^3}{3} + \dots\right)\right) \\ &= \exp\left(n \log(1-x)^{-1}\right) \\ &= \frac{1}{(1-x)^n} \\ &= \sum_{\ell \geq 0} \binom{-n}{\ell} (-x)^\ell \\ &= \sum_{\ell \geq 0} \binom{n+\ell-1}{\ell} x^\ell, \end{aligned}$$

the last step following from the easily checked equality $\binom{-n}{\ell} = (-1)^\ell \binom{n+\ell-1}{\ell}$. Equating coefficients of x^ℓ in the first and last terms of the above string of equalities gives

$$\begin{aligned} Z_{\mathfrak{S}_\ell}(n, n, \dots) &= \binom{n+\ell-1}{\ell} \\ &= \frac{n(n+1)\cdots(n+\ell-1)}{\ell!}, \end{aligned}$$

agreeing with Theorem 7.5 and equation (7.7).

Theorem 7.13 has many enumerative applications. We give one such result here as an example.

7.14 Proposition. *Let $f(n)$ be the number of permutations $\pi \in \mathfrak{S}_n$ of odd order. Equivalently, $\pi^k = \iota$ (the identity permutation) for some odd k . Then*

$$f(n) = \begin{cases} 1^2 \cdot 3^2 \cdot 5^2 \cdots (n-1)^2, & n \text{ even} \\ 1^2 \cdot 3^2 \cdot 5^2 \cdots (n-2)^2 \cdot n, & n \text{ odd.} \end{cases}$$

Proof. A permutation has odd order if and only if all its cycle lengths are odd. Hence [why?]

$$f(n) = n! Z_{\mathfrak{S}_n}(z_i = 1, i \text{ odd}; z_i = 0, i \text{ even}).$$

Making this substitution in Theorem 7.13 gives

$$\sum_{n \geq 0} f(n) \frac{x^n}{n!} = \exp \left(x + \frac{x^3}{3} + \frac{x^5}{5} + \cdots \right).$$

Since $-\log(1-x) = x + \frac{x^2}{2} + \frac{x^3}{3} + \cdots$, we get [why?]

$$\begin{aligned} \sum_{n \geq 0} f(n) \frac{x^n}{n!} &= \exp \left(\frac{1}{2} (-\log(1-x) + \log(1+x)) \right) \\ &= \exp \frac{1}{2} \log \left(\frac{1+x}{1-x} \right) \\ &= \sqrt{\frac{1+x}{1-x}}. \end{aligned}$$

We therefore need to find the coefficients in the power series expansion of $\sqrt{(1+x)/(1-x)}$ at $x = 0$. There is a simple trick for doing so:

$$\begin{aligned} \sqrt{\frac{1+x}{1-x}} &= (1+x)(1-x^2)^{-1/2} \\ &= (1+x) \sum_{m \geq 0} \binom{-1/2}{m} (-x^2)^m \\ &= \sum_{m \geq 0} (-1)^m \binom{-1/2}{m} (x^{2m} + x^{2m+1}), \end{aligned}$$

where by definition

$$\binom{-1/2}{m} = \frac{1}{m!} \left(-\frac{1}{2}\right) \left(-\frac{3}{2}\right) \cdots \left(-\frac{2m-1}{2}\right).$$

It is now a routine computation to check that the coefficient of $x^n/n!$ in $\sqrt{(1+x)/(1-x)}$ agrees with the claimed value of $f(n)$. \square

Quotients of boolean algebras. We will show how to apply Pólya theory to the problem of counting the number of elements of given rank in a quotient poset B_X/G . Here X is a finite set, B_X is the boolean algebra of all subsets of X , and G is a group of permutations of X (with an induced action on B_X). What do colorings of X have to do with subsets? The answer is very simple: a 2-coloring $f: X \rightarrow \{0, 1\}$ corresponds to a subset S_f of X by the usual rule

$$s \in S_f \iff f(s) = 1.$$

Note that two 2-colorings f and g are G -equivalent if and only if S_f and S_g are in the same orbit of G (acting on B_X). Thus the number of inequivalent 2-colorings f of X with i values equal to 1 is just $\#(B_X/G)_i$, the number of elements of B_X/G of rank i . As an immediate application of Pólya's theorem (Theorem 7.7) we obtain the following result.

7.15 Corollary. *We have*

$$\sum_i \#(B_X/G)_i q^i = Z_G(1+q, 1+q^2, 1+q^3, \dots).$$

Proof. If $\kappa(i, j)$ denotes the number of inequivalent 2-colorings of X with the colors 0 and 1 such that 0 is used j times and 1 is used i times (so $i + j = \#X$), then by Pólya's theorem we have

$$\sum_{i,j} \kappa(i, j)x^i y^j = Z_G(x + y, x^2 + y^2, x^3 + y^3, \dots).$$

Setting $x = q$ and $y = 1$ yields the desired result [why?]. \square

Combining Corollary 7.15 with the rank-unimodality of B_X/G (Theorem 5.8) yields the following corollary.

7.16 Corollary. *For any finite group G of permutations of a finite set X , the polynomial $Z_G(1 + q, 1 + q^2, 1 + q^3, \dots)$ has symmetric, unimodal, integer coefficients.*

7.17 Example. (a) For the poset P of Example 5.4(a) we have $G = \{(1)(2)(3), (1, 2)(3)\}$, so $Z_G(z_1, z_2, z_3) = \frac{1}{2}(z_1^3 + z_1 z_2)$. Hence

$$\begin{aligned} \sum_{i=0}^3 (\#P_i)q^i &= \frac{1}{2}((1 + q)^3 + (1 + q)(1 + q^2)) \\ &= 1 + 2q + 2q^2 + q^3. \end{aligned}$$

(b) For the poset P of Example 5.4(b) we have $G = \{(1)(2)(3)(4)(5), (1, 2, 3, 4, 5), (1, 3, 5, 2, 4), (1, 4, 2, 5, 3), (1, 5, 4, 3, 2)\}$, so $Z_G(z_1, z_2, z_3, z_4, z_5) = \frac{1}{5}(z_1^5 + 4z_5)$. Hence

$$\begin{aligned} \sum_{i=0}^5 (\#P_i)q^i &= \frac{1}{5}((1 + q)^5 + 4(1 + q^5)) \\ &= 1 + q + 2q^2 + 2q^3 + q^4 + q^5. \end{aligned}$$

Note that we are equivalently counting 2-colored necklaces (as defined in Example 7.9), say with colors red and blue, of length five according to the number of blue beads. (c) Let X be the squares of a 2×2 chessboard, labelled as follows:

1	2
3	4

Let G be the wreath product $\mathfrak{S}_2 \wr \mathfrak{S}_2$, as defined in Chapter 6. Then

$$G = \{(1)(2)(3)(4), (1, 2)(3)(4), (1)(2)(3, 4), (1, 2)(3, 4), \\ (1, 3)(2, 4), (1, 4)(2, 3), (1, 3, 2, 4), (1, 4, 2, 3)\},$$

so

$$Z_G(z_1, z_2, z_3, z_4) = \frac{1}{8}(z_1^4 + 2z_1^2z_2 + 3z_2^2 + 2z_4).$$

Hence

$$\begin{aligned} \sum_{i=0}^4 (\#P_i)q^i &= \frac{1}{4}((1+q)^4 + 2(1+q)^2(1+q^2) + 3(1+q^2)^2 + 2(1+q^4)) \\ &= 1 + q + 2q^2 + q^3 + q^4 \\ &= \binom{4}{2}, \end{aligned}$$

agreeing with Theorem 6.6.

Using more sophisticated methods (such as the representation theory of the symmetric group), the following generalization of Corollary 7.16 can be proved: let $P(q)$ be any polynomial with symmetric, unimodal, nonnegative, integer coefficients, such as $1 + q + 3q^2 + 3q^3 + 8q^4 + 3q^5 + 3q^6 + q^7 + q^8$ or $q^5 + q^6$ ($= 0 + 0q + \dots + 0q^4 + q^5 + q^6 + 0q^7 + \dots + 0q^{11}$). Then the polynomial $Z_G(P(q), P(q^2), P(q^3), \dots)$ has symmetric, unimodal, nonnegative, integer coefficients.

Graphs. A standard application of Pólya theory is to the enumeration of nonisomorphic graphs. We saw at the end of Chapter 5 that if M is an m -element vertex set, $X = \binom{M}{2}$, and $\mathfrak{S}_m^{(2)}$ is the group of permutations of X induced by permutations of M , then an orbit of i -element subsets of X may be regarded as an isomorphism class of graphs on the vertex set M with i edges. Thus $\#(B_X/\mathfrak{S}_m^{(2)})_i$ is the number of nonisomorphic graphs (without loops or multiple edges) on the vertex set M with i edges. It follows from Corollary 7.15 that if $g_i(m)$ denotes the number of nonisomorphic graphs with m vertices and i edges, then

$$\sum_{i=0}^{\binom{m}{2}} g_i(m)q^i = Z_{\mathfrak{S}_m^{(2)}}(1 + q, 1 + q^2, 1 + q^3, \dots).$$

Thus we would like to compute the cycle enumerator $Z_{\mathfrak{S}_m^{(2)}}(z_1, z_2, \dots)$. If two permutations π and σ of M have the same cycle type (number of cycles of each length), then their actions on X also have the same cycle type [why?]. Thus for each possible cycle type of a permutation of M (i.e., for each partition of m) we need to compute the induced cycle type on X . We also know from Theorem 7.12 the number of permutations of M of each type. For small values of m we can pick some permutation π of each type and compute directly its action on X in order to determine the induced cycle type. For $m = 4$ we have:

CYCLE LENGTHS OF π	NUMBER	π	INDUCED PERMUTATION π'	CYCLE LENGTHS OF π'
1, 1, 1, 1	1	(1)(2)(3)(4)	(12)(13)(14)(23)(24)(34)	1, 1, 1, 1, 1, 1
2, 1, 1	6	(1, 2)(3)(4)	(12)(12, 23)(14, 24)(34)	2, 2, 1, 1
3, 1	8	(1, 2, 3)(4)	(12, 23, 13)(14, 24, 34)	3, 3
2, 2	3	(1, 2)(3, 4)	(12)(13, 24)(14, 23)(34)	2, 2, 1, 1
4	6	(1, 2, 3, 4)	(12, 23, 34, 14)(13, 24)	4, 2

It follows that

$$Z_{\mathfrak{S}_4^{(2)}}(z_1, z_2, z_3, z_4, z_5, z_6) = \frac{1}{24}(z_1^6 + 9z_1^2z_2^2 + 8z_3^2 + 6z_2z_4).$$

If we set $z_i = 1 + q^i$ and simplify, we obtain the polynomial

$$\sum_{i=0}^6 g_i(4)q^i = 1 + q + 2q^2 + 3q^3 + 2q^4 + q^5 + q^6.$$

Indeed, this polynomial agrees with the rank-generating function of the poset of Figure 5.1.

Suppose that we instead want to count the number $h_i(4)$ of nonisomorphic graphs with four vertices and i edges, where now we allow at most *two* edges between any two vertices. We can take M , X , and $G = \mathfrak{S}_4^{(2)}$ as before, but now we have three colors: red for no edges, blue for one edge, and yellow for two edges. A monomial $r^i b^j y^k$ corresponds to a coloring with i pairs of vertices having no edges between them, j pairs having one edge, and k pairs having two edges. The total number e of edges is $j + 2k$. Hence if we let $r = 1, b = q, y = q^2$, then the monomial $r^i b^j y^k$ becomes $q^{j+2k} = q^e$. It follows

that

$$\begin{aligned}
 \sum_{i=0}^{i(i-1)} h_i(4)q^i &= Z_{\mathfrak{S}_4^{(2)}}(1+q+q^2, 1+q^2+q^4, 1+q^3+q^6, \dots) \\
 &= \frac{1}{24} \left((1+q+q^2)^6 + 9(1+q+q^2)^2(1+q^2+q^4)^2 \right. \\
 &\quad \left. + 8(1+q^3+q^6)^2 + 6(1+q^2+q^4)(1+q^4+q^8) \right) \\
 &= 1+q+3q^2+5q^3+8q^4+9q^5+12q^6+9q^7+8q^8+5q^9 \\
 &\quad +3q^{10}+q^{11}+q^{12}.
 \end{aligned}$$

The total number of nonisomorphic graphs on four vertices with edge multiplicities at most two is $\sum_i h_i(4) = 66$.

It should now be clear that if we restrict the edge multiplicity to be r , then the corresponding generating function is $Z_{\mathfrak{S}_4^{(2)}}(1+q+q^2+\dots+q^{r-1}, 1+q^2+q^4+\dots+q^{2r-2}, \dots)$. In particular, to obtain the *total* number $N(r, 4)$ of nonisomorphic graphs on four vertices with edge multiplicity at most r , we simply set each $z_i = r$, obtaining

$$\begin{aligned}
 N(r, 4) &= Z_{\mathfrak{S}_4^{(2)}}(r, r, r, r, r) \\
 &= \frac{1}{24}(r^6 + 9r^4 + 14r^2).
 \end{aligned}$$

This is the same as number of inequivalent r -colorings of the set $X = \binom{M}{2}$ (where $\#M = 4$) [why?].

Of course the same sort of reasoning can be applied to any number of vertices. For five vertices our table becomes the following (using such notation as 1^5 to denote a sequence of five 1's).

CYCLE LENGTHS OF π		NO.	π	INDUCED PERMUTATION	π'	CYCLE LENGTHS OF π'
1^5	1		(1)(2)(3)(4)(5)	(12)(13)...	(45)	1^{10}
$2, 1^3$	10		(1, 2)(3)(4)(5)	(12)(13, 23)(14, 25)(15, 25)(34)(35)(45)		$2^3, 1^4$
$3, 1^2$	20		(1, 2, 3)(4)(5)	(12, 23, 13)(14, 24, 34)(15, 25, 35)(45)		$3^3, 1$
$2^2, 1$	15		(1, 2)(3, 4)(5)	(12)(13, 24)(14, 23)(15, 25)(34)(35, 45)		$2^4, 1^2$
$4, 1$	30		(1, 2, 3, 4)(5)	(12, 23, 34, 14)(13, 24)(15, 25, 35, 45)		$4^2, 2$
$3, 2$	20		(1, 2, 3)(4, 5)	(12, 23, 13)(14, 25, 34, 15, 24, 35)(45)		$6, 3, 1$
5	24		(1, 2, 3, 4, 5)	(12, 23, 34, 45, 15)(13, 24, 35, 14, 25)		5^2

Thus

$$Z_{\mathfrak{S}_5^{(2)}}(z_1, \dots, z_{10}) = \frac{1}{120}(z_1^{10} + 10z_1^4 z_2^3 + 20z_1 z_3^3 + 15z_1^2 z_2^4 + 30z_2 z_4^2 + 20z_1 z_3 z_6 + 24z_5^2),$$

from which we compute

$$\begin{aligned} \sum_{i=0}^{10} g_i(5)q^i &= Z_{\mathfrak{S}_5^{(2)}}(1+q, 1+q^2, \dots, 1+q^{10}) \\ &= 1 + q + 2q^2 + 4q^3 + 6q^4 + 6q^5 + 6q^6 + 4q^7 + 2q^8 + q^9 + q^{10}. \end{aligned}$$

For an arbitrary number $m = \#M$ of vertices there exist explicit formulas for the cycle indicator of the induced action of $\pi \in \mathfrak{S}_M$ on $\binom{M}{2}$, thereby obviating the need to compute π' explicitly as we did in the above tables, but the overall expression for $Z_{\mathfrak{S}_m^{(2)}}$ cannot be simplified significantly or put into a simple generating function as we did in Theorem 7.13. For reference we record

$$\begin{aligned} Z_{\mathfrak{S}_6^{(2)}} &= \frac{1}{6!}(z_1^{15} + 15z_1^7 z_2^4 + 40z_1^3 z_3^4 + 45z_1^3 z_2^6 + 90z_1 z_2 z_4^3 + 120z_1 z_2 z_3^2 z_6 \\ &\quad + 144z_5^3 + 15z_1^3 z_2^6 + 90z_1 z_2 z_4^3 + 40z_3^5 + 120z_3 z_6^2) \end{aligned}$$

$$(g_0(6), g_1(6), \dots, g_{15}(6)) = (1, 1, 2, 5, 9, 15, 21, 24, 24, 21, 15, 9, 5, 2, 1, 1).$$

Moreover if $u(n)$ denotes the number of nonisomorphic simple graphs with n vertices, then

$$(u(0), u(1), \dots, u(11))$$

$$= (1, 1, 2, 4, 11, 34, 156, 1044, 12346, 274668, 12005168, 1018997864).$$

A table of $u(n)$ for $n \leq 75$ is given at

<http://oeis.org/A000088/b000088.txt>

In particular,

$$\begin{aligned}
 u(75) = & 91965776790545918117055311393231179873443957239 \\
 & 0555232344598910500368551136102062542965342147 \\
 & 8723210428876893185920222186100317580740213865 \\
 & 7140377683043095632048495393006440764501648363 \\
 & 4760490012493552274952950606265577383468983364 \\
 & 6883724923654397496226869104105041619919159586 \\
 & 8518775275216748149124234654756641508154401414 \\
 & 8480274454866344981385848105320672784068407907 \\
 & 1134767688676890584660201791139593590722767979 \\
 & 8617445756819562952590259920801220117529208077 \\
 & 0705444809177422214784902579514964768094933848 \\
 & 3173060596932480677345855848701061537676603425 \\
 & 1254842843718829212212327337499413913712750831 \\
 & 0550986833980707875560051306072520155744624852 \\
 & 0263616216031346723897074759199703968653839368 \\
 & 77636080643275926566803872596099072,
 \end{aligned}$$

a number of 726 digits! Compare

$$\frac{2^{\binom{75}{2}}}{75!} = .9196577679054591809 \times 10^{726},$$

which agrees with $u(75)$ to 17 significant digits [why?].

Notes for Chapter 7

Burnside's lemma (Lemma 7.2) was actually first stated and proved by Frobenius [41, end of §4]. Frobenius in turn credits Cauchy [22, p. 286] for proving the lemma in the transitive case. Burnside, in the first edition of his book [18, §118–119], attributes the lemma to Frobenius, but in the second edition [19] this citation is absent. For more on the history of Burnside's lemma, see [81] and [121]. Many authors now call this result the Cauchy-Frobenius lemma. The cycle indicator $Z_G(z_1, z_2, \dots)$ (where G is a subgroup

of \mathfrak{S}_n) was first considered by J.H. Redfield [93], who called it the *group reduction function*, denoted $\text{Grf}(G)$. G. Pólya [86] independently defined the cycle indicator, proved the fundamental Theorem 7.7, and gave numerous applications. For an English translation of Pólya's paper, see [87]. Much of Pólya's work was anticipated by Redfield. For interesting historical information about the work of Redfield and its relation to Pólya theory, see [50][52][71][94] (all in the same issue of *Journal of Graph Theory*). The Wikipedia article "John Howard Redfield" also gives information and references on the interesting story of the rediscovery and significance of Redfield's work.

The application of Pólya's theorem to the enumeration of nonisomorphic graphs appears in Pólya's original paper [86]. For much additional work on graphical enumeration, see the text of Harary and Palmer [51].


Subsequent to Pólya's work there have been a huge number of expositions, applications, and generalizations of Pólya theory. An example of such a generalization appears in Exercise 7.11. We mention here only the nice survey [16] by N. G. de Bruijn.

Theorem 7.13 (the generating function for the cycle indicator $Z_{\mathfrak{S}_\ell}$ of the symmetric group \mathfrak{S}_ℓ) goes back to Frobenius (see [42, bottom of p. 152 of GA]) and Hurwitz [61, §4]. It is clear that they were aware of Theorem 7.13, even if they did not state it explicitly. For a more conceptual approach and further aspects see R. Stanley [110, §§5.1–5.2].

Chapter 8

A glimpse of Young tableaux

We defined in Chapter 6 Young's lattice Y , the poset of all partitions of all nonnegative integers, ordered by containment of their Young diagrams.




Here we will be concerned with the counting of certain walks in the Hasse diagram (considered as a graph) of Y . Note that since Y is infinite, we cannot talk about its eigenvalues and eigenvectors. We need different techniques for counting walks. It will be convenient to denote the length of a walk by n , rather than by ℓ as in previous chapters.

Note that Y is a graded poset (of infinite rank), with Y_i consisting of all partitions of i . In other words, we have $Y = Y_0 \cup Y_1 \cup \dots$ (disjoint union),

where every maximal chain intersects each level Y_i exactly once. We call Y_i the i th *level* of Y , just as we did for finite graded posets.

Since the Hasse diagram of Y is a simple graph (no loops or multiple edges), a walk of length n is specified by a sequence $\lambda^0, \lambda^1, \dots, \lambda^n$ of vertices of Y . We will call a walk in the Hasse diagram of a poset a *Hasse walk*. Each λ^i is a partition of some integer, and we have either (a) $\lambda^i < \lambda^{i+1}$ and $|\lambda^i| = |\lambda^{i+1}| - 1$, or (b) $\lambda^i > \lambda^{i+1}$ and $|\lambda^i| = |\lambda^{i+1}| + 1$. (Recall that for a partition λ , we write $|\lambda|$ for the sum of the parts of λ .) A step of type (a) is denoted by U (for “up,” since we move up in the Hasse diagram), while a step of type (b) is denoted by D (for “down”). If the walk W has steps of types A_1, A_2, \dots, A_n , respectively, where each A_i is either U or D , then we say that W is of *type* $A_n A_{n-1} \cdots A_2 A_1$. Note that the type of a walk is written in the *opposite* order to that of the walk. This is because we will soon regard U and D as linear transformations, and we multiply linear transformations *right-to-left* (opposite to the usual left-to-right reading order). For instance (abbreviating a partition $(\lambda_1, \dots, \lambda_m)$ as $\lambda_1 \cdots \lambda_m$), the walk $\emptyset, 1, 2, 1, 11, 111, 211, 221, 22, 21, 31, 41$ is of type $UUDDUUUUUU = U^2 D^2 U^4 D U^2$.

There is a nice combinatorial interpretation of walks of type U^n which begin at \emptyset . Such walks are of course just saturated chains $\emptyset = \lambda^0 < \lambda^1 < \cdots < \lambda^n$. In other words, they may be regarded as sequences of Young diagrams, beginning with the empty diagram and adding one new square at each step. An example of a walk of type U^5 is given by



We can specify this walk by taking the final diagram and inserting an i into square s if s was added at the i th step. Thus the above walk is encoded by the “tableau”

1	2
3	5
4	

.

Such an object τ is called a *standard Young tableaux* (or SYT). It consists of the Young diagram D of some partition λ of an integer n , together with the numbers $1, 2, \dots, n$ inserted into the squares of D , so that each number appears exactly once, and every row and column is *increasing*. We call λ the *shape* of the SYT τ , denoted $\lambda = \text{sh}(\tau)$. For instance, there are five SYT of shape $(2, 2, 1)$, given by

1	2	1	2	1	3	1	3	1	4
3	4	3	5	2	4	2	5	2	5
5		4		5		4		3	

Let f^λ denote the number of SYT of shape λ , so for instance $f^{(2,2,1)} = 5$. The numbers f^λ have many interesting properties; for instance, there is a famous explicit formula for them known as the Frame–Robinson–Thrall hook length formula. For the sake of completeness we state this formula without proof, though it is not needed in what follows.

Let u be a square of the Young diagram of the partition λ . Define the *hook* $H(u)$ of u (or at u) to be the set of all squares directly to the right of u or directly below u , including u itself. The size (number of squares) of $H(u)$ is called the *hook length* of u (or at u), denoted $h(u)$. In the diagram of the partition $(4, 2, 2)$ below, we have inserted the hook length $h(u)$ inside each square u .

6	5	2	1
3	2		
2	1		

8.1 Theorem (hook length formula). *Let $\lambda \vdash n$. Then*

$$f^\lambda = \frac{n!}{\prod_{u \in \lambda} h(u)}.$$

Here the notation $u \in \lambda$ means that u ranges over all squares of the Young diagram of λ .

For instance, the diagram of the hook lengths of $\lambda = (4, 2, 2)$ above gives

$$f^{(4,2,2)} = \frac{8!}{6 \cdot 5 \cdot 2 \cdot 1 \cdot 3 \cdot 2 \cdot 2 \cdot 1} = 56.$$

In this chapter we will be concerned with the connection between SYT and counting walks in Young's lattice. If $w = A_n A_{n-1} \cdots A_1$ is some word in U and D and $\lambda \vdash n$, then let us write $\alpha(w, \lambda)$ for the number of Hasse walks in Y of type w which start at the empty partition \emptyset and end at λ . For instance, $\alpha(UDUU, 11) = 2$, the corresponding walks being $\emptyset, 1, 2, 1, 11$ and $\emptyset, 1, 11, 1, 11$. Thus in particular $\alpha(U^n, \lambda) = f^\lambda$ [why?]. In a similar fashion, since the number of Hasse walks of type $D^n U^n$ which begin at \emptyset , go up to a partition $\lambda \vdash n$, and then back down to \emptyset is given by $(f^\lambda)^2$, we have

$$\alpha(D^n U^n, \emptyset) = \sum_{\lambda \vdash n} (f^\lambda)^2. \quad (8.1)$$

Our object is to find an explicit formula for $\alpha(w, \lambda)$ of the form $f^\lambda c_w$, where c_w does not depend on λ . (It is by no means *a priori* obvious that such a formula should exist.) In particular, since $f^\emptyset = 1$, we will obtain by setting $\lambda = \emptyset$ a simple formula for the number of (closed) Hasse walks of type w from \emptyset to \emptyset (thus including a simple formula for (8.1)).

There is an easy condition for the existence of *any* Hasse walk of type w from \emptyset to λ , given by the next lemma.

8.2 Lemma. *Suppose $w = D^{s_k} U^{r_k} \cdots D^{s_2} U^{r_2} D^{s_1} U^{r_1}$, where $r_i \geq 0$ and $s_i \geq 0$. Let $\lambda \vdash n$. Then there exists a Hasse walk of type w from \emptyset to λ if and only if:*

$$\sum_{i=1}^k (r_i - s_i) = n$$

$$\sum_{i=1}^j (r_i - s_i) \geq 0 \text{ for } 1 \leq j \leq k.$$

Proof. Since each U moves up one level and each D moves down one level, we see that $\sum_{i=1}^k (r_i - s_i)$ is the level at which a walk of type w beginning at \emptyset ends. Hence $\sum_{i=1}^k (r_i - s_i) = |\lambda| = n$.

After $\sum_{i=1}^j (r_i + s_i)$ steps we will be at level $\sum_{i=1}^j (r_i - s_i)$. Since the lowest level is level 0, we must have $\sum_{i=1}^j (r_i - s_i) \geq 0$ for $1 \leq j \leq k$.

The easy proof that the two conditions of the lemma are *sufficient* for the existence of a Hasse walk of type w from \emptyset to λ is left to the reader. \square

If w is a word in U and D satisfying the conditions of Lemma 8.2, then we say that w is a *valid λ -word*. Note that the condition of being a valid λ -word depends only on $|\lambda|$.

The proof of our formula for $\alpha(w, \lambda)$ will be based on linear transformations analogous to those defined by equation (4.2) and (4.3). As in Chapter 4 let $\mathbb{R}Y_j$ be the real vector space with basis Y_j . Define two linear transformations $U_i: \mathbb{R}Y_i \rightarrow \mathbb{R}Y_{i+1}$ and $D_i: \mathbb{R}Y_i \rightarrow \mathbb{R}Y_{i-1}$ by

$$U_i(\lambda) = \sum_{\substack{\mu \vdash i+1 \\ \lambda < \mu}} \mu$$

$$D_i(\lambda) = \sum_{\substack{\nu \vdash i-1 \\ \nu < \lambda}} \nu,$$

for all $\lambda \vdash i$. For instance (using abbreviated notation for partitions)

$$U_{21}(54422211) = 64422211 + 55422211 + 54432211 + 54422221 + 544222111$$

$$D_{21}(54422211) = 44422211 + 54322211 + 54422111 + 5442221.$$

It is clear [why?] that if r is the number of *distinct* (i.e., unequal) parts of λ , then $U_i(\lambda)$ is a sum of $r + 1$ terms and $D_i(\lambda)$ is a sum of r terms. The next lemma is an analogue for Y of the corresponding result for B_n (Lemma 4.6).

8.3 Lemma. *For any $i \geq 0$ we have*

$$D_{i+1}U_i - U_{i-1}D_i = I_i, \tag{8.2}$$

the identity linear transformation on $\mathbb{R}Y_i$.

Proof. Apply the left-hand side of (8.2) to a partition λ of i , expand in terms of the basis Y_i , and consider the coefficient of a partition μ . If $\mu \neq \lambda$ and μ can be obtained from λ by adding one square s to (the Young diagram of) λ and then removing a (necessarily different) square t , then there is exactly one choice of s and t . Hence the coefficient of μ in $D_{i+1}U_i(\lambda)$ is equal to 1. But then there is exactly one way to remove a square from λ and then add a square to get μ , namely, remove t and add s . Hence the coefficient of μ in $U_{i-1}D_i(\lambda)$ is also 1, so the coefficient of μ when the left-hand side of (8.2) is applied to λ is 0.

If now $\mu \neq \lambda$ and we cannot obtain μ by adding a square and then deleting a square from λ (i.e., μ and λ differ in more than two rows), then clearly when we apply the left-hand side of (8.2) to λ , the coefficient of μ will be 0.

Finally consider the case $\lambda = \mu$. Let r be the number of distinct (unequal) parts of λ . Then the coefficient of λ in $D_{i+1}U_i(\lambda)$ is $r+1$, while the coefficient of λ in $U_{i-1}D_i(\lambda)$ is r , since there are $r+1$ ways to add a square to λ and then remove it, while there are r ways to remove a square and then add it back in. Hence when we apply the left-hand side of (8.2) to λ , the coefficient of λ is equal to 1.

Combining the conclusions of the three cases just considered shows that the left-hand side of (8.2) is just I_i , as was to be proved. \square

We come to one of the main results of this chapter.

8.4 Theorem. *Let λ be a partition and $w = A_n A_{n-1} \cdots A_1$ a valid λ -word. Let $S_w = \{i : A_i = D\}$. For each $i \in S_w$, let a_i be the number of D 's in w to the right of A_i , and let b_i be the number of U 's in w to the right of A_i . Thus $a_i - b_i$ is the level we occupy in Y before taking the step $A_i = D$. Then*

$$\alpha(w, \lambda) = f^\lambda \prod_{i \in S_w} (b_i - a_i).$$

Before proving Theorem 8.4, let us give an example. Suppose $w = U^3 D^2 U^2 D U^3 = U U U D D U U D U U U$ and $\lambda = (2, 2, 1)$. Then $S_w = \{4, 7, 8\}$ and $a_4 = 0$, $b_4 = 3$, $a_7 = 1$, $b_7 = 5$, $a_8 = 2$, $b_8 = 5$. We have also seen earlier that $f^{221} = 5$. Thus

$$\alpha(w, \lambda) = 5(3 - 0)(5 - 1)(5 - 2) = 180.$$

Proof. Proof of Theorem 8.4. For notational simplicity we will omit the subscripts from the linear transformations U_i and D_i . This should cause no confusion since the subscripts will be uniquely determined by the elements on which U and D act. For instance, the expression $UDUU(\lambda)$ where $\lambda \vdash i$ must mean $U_{i+1}D_{i+2}U_{i+1}U_i(\lambda)$; otherwise it would be undefined since U_j and D_j can only act on elements of $\mathbb{R}Y_j$, and moreover U_j raises the level by one while D_j lowers it by one.

By (8.2) we can replace DU in any word y in the letters U and D by $UD + I$. This replaces y by a sum of two words, one with one fewer D and the other with one D moved one space to the right. For instance, replacing the first DU in $UUDUDDU$ by $UD + I$ yields $UUUDDDU + UUDDU$. If

we begin with the word w and iterate this procedure, replacing a DU in any word with $UD + I$, eventually there will be no U 's to the right of any D 's and the procedure will come to an end. At this point we will have expressed w as a linear combination (with integer coefficients) of words of the form $U^i D^j$. Since the operation of replacing DU with $UD + I$ preserves the difference between the number of U 's and D 's in each word, all the words $U^i D^j$ which appear will have $i - j$ equal to some constant n (namely, the number of U 's minus the number of D 's in w). Specifically, say we have

$$w = \sum_{i-j=n} r_{ij}(w) U^i D^j, \quad (8.3)$$

where each $r_{ij}(w) \in \mathbb{Z}$. (We also define $r_{ij}(w) = 0$ if $i < 0$ or $j < 0$.) We claim that the $r_{ij}(w)$'s are uniquely determined by w . Equivalently [why?], if we have

$$\sum_{i-j=n} d_{ij} U^i D^j = 0 \quad (8.4)$$

(as an identity of linear transformations acting on the space $\mathbb{R}Y_k$ for *any* k), where each $d_{ij} \in \mathbb{Z}$ (or $d_{ij} \in \mathbb{R}$, if you prefer), then each $d_{ij} = 0$. Let j' be the least integer for which $d_{j'+n, j'} \neq 0$. Let $\mu \vdash j'$, and apply both sides of equation (8.4) to μ . The left-hand side has exactly one nonzero term, namely, the term with $j = j'$ [why?]. The right-hand side, on the other hand¹, is 0, a contradiction. Thus the $r_{ij}(w)$'s are unique.

Now apply U on the left to equation (8.3). We get

$$Uw = \sum_{i,j} r_{ij}(w) U^{i+1} D^j.$$

Hence (using uniqueness of the r_{ij} 's) there follows [why?]

$$r_{ij}(Uw) = r_{i-1, j}(w). \quad (8.5)$$

We next want to apply D on the left to (8.3). It is easily proved by induction on i (left as an exercise) that

$$DU^i = U^i D + iU^{i-1}. \quad (8.6)$$

¹The phrase "the right-hand side, on the other hand" does not mean the left-hand side!

(We interpret U^{-1} as being 0, so that equation (8.6) is true for $i = 0$.) Hence

$$\begin{aligned} Dw &= \sum_{i,j} r_{ij}(w) DU^i D^j \\ &= \sum_{i,j} r_{ij}(w) (U^i D + iU^{i-1}) D^j, \end{aligned}$$

from which it follows [why?] that

$$r_{ij}(Dw) = r_{i,j-1}(w) + (i+1)r_{i+1,j}(w). \quad (8.7)$$

Setting $j = 0$ in (8.5) and (8.7) yields

$$r_{i0}(Uw) = r_{i-1,0}(w) \quad (8.8)$$

$$r_{i0}(Dw) = (i+1)r_{i+1,0}(w). \quad (8.9)$$

Now let (8.3) operate on \emptyset . Since $D^j(\emptyset) = 0$ for all $j > 0$, we get $w(\emptyset) = r_{n0}(w)U^n(\emptyset)$. Thus the coefficient of λ in $w(\emptyset)$ is given by

$$\alpha(w, \lambda) = r_{n0}(w)\alpha(U^n, \lambda) = r_{n0}f^\lambda,$$

where as usual $\lambda \vdash n$. It is clear from (8.8) and (8.9) that

$$r_{n0}(w) = \prod_{j \in S_w} (b_j - a_j),$$

and the proof follows. □

NOTE. It is possible to give a simpler proof of Theorem 8.4, but the proof we have given is useful for generalizations not appearing here.

An interesting special case of the previous theorem allows us to evaluate equation (8.1).

8.5 Corollary. *We have*

$$\alpha(D^n U^n, \emptyset) = \sum_{\lambda \vdash n} (f^\lambda)^2 = n!.$$

Proof. When $w = D^n U^n$ in Theorem 8.4 we have $S_w = \{n + 1, n + 2, \dots, 2n\}$, $a_i = n - i + 1$, and $b_i = n$, from which the proof is immediate. \square

NOTE (for those familiar with the representation theory of finite groups). It can be shown that the numbers f^λ , for $\lambda \vdash n$, are the degrees of the irreducible representations of the symmetric group \mathfrak{S}_n . Given this, Corollary 8.5 is a special case of the result that the sum of the squares of the degrees of the irreducible representations of a finite group G is equal to the order $\#G$ of G . There are many other intimate connections between the representation theory of \mathfrak{S}_n , on the one hand, and the combinatorics of Young's lattice and Young tableaux, on the other. There is also an elegant combinatorial proof of Corollary 8.5, based on the *RSK algorithm* (after Gilbert de Beauregard Robinson, Craige Schensted, and Donald Knuth) or *Robinson–Schensted correspondence*, with many fascinating properties and with deep connections to representation theory. In the first Appendix at the end of this chapter we give a description of the RSK algorithm and the combinatorial proof of Corollary 8.5.

We now consider a variation of Theorem 8.4 in which we are not concerned with the type w of a Hasse walk from \emptyset to w , but only with the number of steps. For instance, there are three Hasse walks of length three from \emptyset to the partition 1, given by $\emptyset, 1, \emptyset, 1$; $\emptyset, 1, 2, 1$; and $\emptyset, 1, 11, 1$. Let $\beta(\ell, \lambda)$ denote the number of Hasse walks of length ℓ from \emptyset to λ . Note the two following easy facts:

(F1) $\beta(\ell, \lambda) = 0$ unless $\ell \equiv |\lambda| \pmod{2}$.

(F2) $\beta(\ell, \lambda)$ is the coefficient of λ in the expansion of $(D + U)^\ell(\emptyset)$ as a linear combination of partitions.

Because of (F2) it is important to write $(D + U)^\ell$ as a linear combination of terms $U^i D^j$, just as in the proof of Theorem 8.4 we wrote a word w in U and D in this form. Thus define integers $b_{ij}(\ell)$ by

$$(D + U)^\ell = \sum_{i,j} b_{ij}(\ell) U^i D^j. \quad (8.10)$$

Just as in the proof of Theorem 8.4, the numbers $b_{ij}(\ell)$ exist and are well-defined.

8.6 Lemma. *We have $b_{ij}(\ell) = 0$ if $\ell - i - j$ is odd. If $\ell - i - j = 2m$ then*

$$b_{ij}(\ell) = \frac{\ell!}{2^m i! j! m!}. \quad (8.11)$$

Proof. The assertion for $\ell - i - j$ odd is equivalent to (F1) above, so assume $\ell - i - j$ is even. The proof is by induction on ℓ . It's easy to check that (8.11) holds for $\ell = 1$. Now assume true for some fixed $\ell \geq 1$. Using (8.10) we obtain

$$\begin{aligned} \sum_{i,j} b_{ij}(\ell + 1)U^i D^j &= (D + U)^{\ell+1} \\ &= (D + U) \sum_{i,j} b_{ij}(\ell)U^i D^j \\ &= \sum_{i,j} b_{ij}(\ell)(DU^i D^j + U^{i+1} D^j). \end{aligned}$$

In the proof of Theorem 8.4 we saw that $DU^i = U^i D + iU^{i-1}$ (see equation (8.6)). Hence we get

$$\sum_{i,j} b_{ij}(\ell + 1)U^i D^j = \sum_{i,j} b_{ij}(\ell)(U^i D^{j+1} + iU^{i-1} D^j + U^{i+1} D^j). \quad (8.12)$$

As mentioned after (8.10), the expansion of $(D + U)^{\ell+1}$ in terms of $U^i D^j$ is unique. Hence equating coefficients of $U^i D^j$ on both sides of (8.12) yields the recurrence

$$b_{ij}(\ell + 1) = b_{i,j-1}(\ell) + (i + 1)b_{i+1,j}(\ell) + b_{i-1,j}(\ell). \quad (8.13)$$

It is a routine matter to check that the function $\ell!/2^m i! j! m!$ satisfies the same recurrence (8.13) as $b_{ij}(\ell)$, with the same initial condition $b_{00}(0) = 1$. From this the proof follows by induction. \square

From Lemma 8.6 it is easy to prove the following result.

8.7 Theorem. *Let $\ell \geq n$ and $\lambda \vdash n$, with $\ell - n$ even. Then*

$$\beta(\ell, \lambda) = \binom{\ell}{n} (1 \cdot 3 \cdot 5 \cdots (\ell - n - 1)) f^\lambda.$$

Proof. Apply both sides of (8.10) to \emptyset . Since $U^i D^j(\emptyset) = 0$ unless $j = 0$, we get

$$\begin{aligned} (D + U)^\ell(\emptyset) &= \sum_i b_{i0}(\ell)U^i(\emptyset) \\ &= \sum_i b_{i0}(\ell) \sum_{\lambda \vdash i} f^\lambda. \end{aligned}$$

Since by Lemma 8.6 we have $b_{i0}(\ell) = \binom{\ell}{i}(1 \cdot 3 \cdot 5 \cdots (\ell - i - 1))$ when $\ell - i$ is even, the proof follows from (F2). \square

NOTE. The proof of Theorem 8.7 only required knowing the value of $b_{i0}(\ell)$. However, in Lemma 8.6 we computed $b_{ij}(\ell)$ for all j . We could have carried out the proof so as only to compute $b_{i0}(\ell)$, but the general value of $b_{ij}(\ell)$ is so simple that we have included it too.

8.8 Corollary. *The total number of Hasse walks in Y of length $2m$ from \emptyset to \emptyset is given by*

$$\beta(2m, \emptyset) = 1 \cdot 3 \cdot 5 \cdots (2m - 1).$$

Proof. Simply substitute $\lambda = \emptyset$ (so $n = 0$) and $\ell = 2m$ in Theorem 8.7. \square

The fact that we can count various kinds of Hasse walks in Y suggests that there may be some finite graphs related to Y whose eigenvalues we can also compute. This is indeed the case, and we will discuss the simplest case here. (See Exercise 8.21 for a generalization.) Let $Y_{j-1,j}$ denote the restriction of Young's lattice Y to ranks $j - 1$ and j . Identify $Y_{j-1,j}$ with its Hasse diagram, regarded as a (bipartite) graph. Let $p(i) = \#Y_i$, the number of partitions of i .

8.9 Theorem. *The eigenvalues of $Y_{j-1,j}$ are given as follows: 0 is an eigenvalue of multiplicity $p(j) - p(j - 1)$; and for $1 \leq s \leq j$, the numbers $\pm\sqrt{s}$ are eigenvalues of multiplicity $p(j - s) - p(j - s - 1)$.*

Proof. Let \mathbf{A} denote the adjacency matrix of $Y_{j-1,j}$. Since $\mathbb{R}Y_{j-1,j} = \mathbb{R}Y_{j-1} \oplus \mathbb{R}Y_j$ (vector space direct sum), any vector $v \in \mathbb{R}Y_{j-1,j}$ can be written uniquely as $v = v_{j-1} + v_j$, where $v_i \in \mathbb{R}Y_i$. The matrix \mathbf{A} acts on the vector space $\mathbb{R}Y_{j-1,j}$ as follows [why?]:

$$\mathbf{A}(v) = D(v_j) + U(v_{j-1}). \tag{8.14}$$

Just as Theorem 4.7 followed from Lemma 4.6, we deduce from Lemma 8.3 that for any i we have that $U_i: \mathbb{R}Y_i \rightarrow \mathbb{R}Y_{i+1}$ is one-to-one and $D_i: \mathbb{R}Y_i \rightarrow \mathbb{R}Y_{i-1}$ is onto. It follows in particular that

$$\begin{aligned} \dim \ker(D_i) &= \dim \mathbb{R}Y_i - \dim \mathbb{R}Y_{i-1} \\ &= p(i) - p(i - 1), \end{aligned}$$

where \ker denotes kernel.

Case 1. Let $v \in \ker(D_j)$, so $v = v_j$. Then $\mathbf{A}v = Dv = 0$. Thus $\ker(D_j)$ is an eigenspace of \mathbf{A} for the eigenvalue 0, so 0 is an eigenvalue of multiplicity at least $p(j) - p(j-1)$.

Case 2. Let $v \in \ker(D_s)$ for some $0 \leq s \leq j-1$. Let

$$v^* = \pm\sqrt{j-s}U^{j-1-s}(v) + U^{j-s}(v).$$

Note that $v^* \in \mathbb{R}Y_{j-1,j}$, with $v_{j-1}^* = \pm\sqrt{j-s}U^{j-1-s}(v)$ and $v_j^* = U^{j-s}(v)$. Using equation (8.6), we compute

$$\begin{aligned} \mathbf{A}(v^*) &= U(v_{j-1}^*) + D(v_j^*) \\ &= \pm\sqrt{j-s}U^{j-s}(v) + DU^{j-s}(v) \\ &= \pm\sqrt{j-s}U^{j-s}(v) + U^{j-s}D(v) + (j-s)U^{j-s-1}(v) \\ &= \pm\sqrt{j-s}U^{j-s}(v) + (j-s)U^{j-s-1}(v) \\ &= \pm\sqrt{j-s}v^*. \end{aligned} \tag{8.15}$$

It's easy to verify (using the fact that U is one-to-one) that if $v(1), \dots, v(t)$ is a basis for $\ker(D_s)$, then $v(1)^*, \dots, v(t)^*$ are linearly independent. Hence by (8.15) we have that $\pm\sqrt{j-s}$ is an eigenvalue of \mathbf{A} of multiplicity at least $t = \dim \ker(D_s) = p(s) - p(s-1)$.

We have found a total of

$$p(j) - p(j-1) + 2 \sum_{s=0}^{j-1} (p(s) - p(s-1)) = p(j-1) + p(j)$$

eigenvalues of \mathbf{A} . (The factor 2 above arises from the fact that both $+\sqrt{j-s}$ and $-\sqrt{j-s}$ are eigenvalues.) Since the graph $Y_{j-1,j}$ has $p(j-1) + p(j)$ vertices, we have found all its eigenvalues. \square

An elegant combinatorial consequence of Theorem 8.9 is the following.

8.10 Corollary. *Fix $j \geq 1$. The number of ways to choose a partition λ of j , then delete a square from λ (keeping it a partition), then insert a square, then delete a square, etc., for a total of m insertions and m deletions, ending back at λ , is given by*

$$\sum_{s=1}^j [p(j-s) - p(j-s-1)]s^m, \quad m > 0. \tag{8.16}$$

Proof. Exactly half the closed walks in $Y_{j-1,j}$ of length $2m$ begin at an element of Y_j [why?]. Hence if $Y_{j-1,j}$ has eigenvalues $\theta_1, \dots, \theta_r$, then by Corollary 1.3 the desired number of walks is given by $\frac{1}{2}(\theta_1^{2m} + \dots + \theta_r^{2m})$. Using the values of $\theta_1, \dots, \theta_r$ given by Theorem 8.9 yields (8.16). \square

For instance, when $j = 7$, equation (8.16) becomes $4 + 2 \cdot 2^m + 2 \cdot 3^m + 4^m + 5^m + 7^m$. When $m = 1$ we get 30, the number of edges of the graph $Y_{6,7}$ [why?].

APPENDIX 1: THE RSK ALGORITHM

We will describe a bijection between permutations $\pi \in \mathfrak{S}_n$ and pairs (P, Q) of SYT of the same shape $\lambda \vdash n$. Define a *near Young tableau* (NYT) to be the same as an SYT, except that the entries can be any distinct integers, not necessarily the integers $1, 2, \dots, n$. Let P_{ij} denote the entry in row i and column j of P . The basic operation of the RSK algorithm consists of the *row insertion* $P \leftarrow k$ of a positive integer k into an NYT $P = (P_{ij})$. The operation $P \leftarrow k$ is defined as follows: let r be the least integer such that $P_{1r} > k$. If no such r exists (i.e., all elements of the first row of P are less than k), then simply place k at the end of the first row. The insertion process stops, and the resulting NYT is $P \leftarrow k$. If, on the other hand, r does exist then replace P_{1r} by k . The element k then “bumps” $P_{1r} := k'$ into the second row, i.e., insert k' into the second row of P by the insertion rule just described. Either k' is inserted at the end of the second row, or else it bumps an element k'' to the third row. Continue until an element is inserted at the end of a row (possibly as the first element of a new row). The resulting array is $P \leftarrow k$.

8.11 Example. Let

$$P = \begin{array}{cccc} & 3 & 7 & 9 & 14 \\ & & 6 & 11 & 12 \\ 10 & & & & 16 \\ & & & 13 & \\ & & & & 15 \end{array}$$

Then $P \leftarrow 8$ is shown below, with the elements inserted into each row (either by bumping or by the final insertion in the fourth row) in boldface. Thus the 8 bumps the 9, the 9 bumps the 11, the 11 bumps the 16, and the 16 is inserted at the end of a row. Hence

$$(P \leftarrow 8) = \begin{array}{cccc} & 3 & 7 & \mathbf{8} & 14 \\ & & 6 & \mathbf{9} & 12 \\ 10 & & & \mathbf{11} & \\ & & & 13 & \mathbf{16} \\ & & & & 15 \end{array} .$$

We omit the proof, which is fairly straightforward, that if P is an NYT, then so is $P \leftarrow k$. We can now describe the RSK algorithm. Let $\pi =$

$a_1 a_2 \cdots a_n \in \mathfrak{S}_n$. We will inductively construct a sequence $(P_0, Q_0), (P_1, Q_1), \dots, (P_n, Q_n)$ of pairs (P_i, Q_i) of NYT of the same shape, where P_i and Q_i each have i squares. First, define $(P_0, Q_0) = (\emptyset, \emptyset)$. If (P_{i-1}, Q_{i-1}) have been defined, then set $P_i = P_{i-1} \leftarrow a_i$. In other words, P_i is obtained from P_{i-1} by row inserting a_i . Now define Q_i to be the NYT obtained from Q_{i-1} by inserting i so that Q_i and P_i have the same shape. (The entries of Q_{i-1} don't change; we are simply placing i into a certain new square and not row-inserting it into Q_{i-1} .) Finally let $(P, Q) = (P_n, Q_n)$. We write $\pi \xrightarrow{\text{RSK}} (P, Q)$.

8.12 Example. Let $\pi = 4273615 \in \mathfrak{S}_7$. The pairs $(P_1, Q_1), \dots, (P_7, Q_7) = (P, Q)$ are as follows:

<u>P_i</u>	<u>Q_i</u>
4	1
2 4	1 2
2 7 4	1 3 2
2 3 4 7	1 3 2 4
2 3 6 4 7	1 3 5 2 4
1 3 6 2 7 4	1 3 5 2 4 6
1 3 5 2 6 4 7	1 3 5 2 4 6 7

8.13 Theorem. *The RSK algorithm defines a bijection between the symmetric group \mathfrak{S}_n and the set of all pairs (P, Q) of SYT of the same shape, where the shape λ is a partition of n .*

Proof (sketch). The key step is to define the inverse of RSK. In other words, if $\pi \mapsto (P, Q)$, then how can we recover π uniquely from (P, Q) ? Moreover, we need to find π for *any* (P, Q) . Observe that the position occupied by

n in Q is the last position to be occupied in the insertion process. Suppose that k occupies this position in P . It was bumped into this position by some element j in the row above k that is currently the largest element of its row less than k . Hence we can “inverse bump” k into the position occupied by j , and now inverse bump j into the row above it by the same procedure. Eventually an element will be placed in the first row, inverse bumping another element t out of the tableau altogether. Thus t was the last element of π to be inserted, i.e., if $\pi = a_1 a_2 \cdots a_n$ then $a_n = t$. Now locate the position occupied by $n-1$ in Q and repeat the procedure, obtaining a_{n-1} . Continuing in this way, we uniquely construct π one element at a time from right-to-left, such that $\pi \mapsto (P, Q)$. \square

The RSK-algorithm provides a bijective proof of Corollary 8.5, that is,

$$\sum_{\lambda \vdash n} (f^\lambda)^2 = n!.$$

APPENDIX 2: PLANE PARTITIONS

In this appendix we show how a generalization of the RSK algorithm leads to an elegant generating function for a two-dimensional generalization of integer partitions. A *plane partition* of an integer $n \geq 0$ is a two-dimensional array $\pi = (\pi_{ij})_{i,j \geq 1}$ of integers $\pi_{ij} \geq 0$ that is weakly decreasing in rows and columns, i.e.,

$$\pi_{ij} \geq \pi_{i+1,j}, \quad \pi_{ij} \geq \pi_{i,j+1},$$

such that $\sum_{i,j} \pi_{ij} = n$. It follows that all but finitely many π_{ij} are 0, and these 0's are omitted in writing a particular plane partition π . Given a plane partition π , we write $|\pi| = n$ to denote that π is a plane partition of n . More generally, if L is any array of nonnegative integers we write $|L|$ for the sum of the parts (entries) of L .

There is one plane partition of 0, namely, all $\pi_{ij} = 0$, denoted \emptyset . The plane partitions of the integers $0 \leq n \leq 3$ are given by

$$\begin{array}{cccccccccccc} \emptyset & 1 & 2 & 11 & 1 & 3 & 21 & 111 & 11 & 2 & 1 \\ & & & & 1 & & & & 1 & 1 & 1 \\ & & & & & & & & & & 1 \end{array} .$$

If $\text{pp}(n)$ denotes the number of plane partitions of n , then $\text{pp}(0) = 1$, $\text{pp}(1) = 1$, $\text{pp}(2) = 3$, and $\text{pp}(3) = 6$.

Our object is to give a formula for the generating function

$$F(x) = \sum_{n \geq 0} \text{pp}(n)x^n = 1 + x + 3x^2 + 6x^3 + 13x^4 + 24x^5 + \cdots .$$

More generally, we will consider plane partitions with at most r rows and at most s columns, i.e., $\pi_{ij} = 0$ for $i > r$ or $j > s$. As a simple warmup, let us first consider the case of ordinary partitions $\lambda = (\lambda_1, \lambda_2, \dots)$ of n .

8.14 Proposition. *Let $p_s(n)$ denote the number of partitions of n with at most s parts. Equivalently, $p_s(n)$ is the number of plane partitions of n with at most one row and at most s columns [why?]. Then*

$$\sum_{n \geq 0} p_s(n)x^n = \prod_{k=1}^s (1 - x^k)^{-1}.$$

Proof. First note that the partition λ has at most s parts if and only if the conjugate partition λ' defined in Chapter 6 has largest part at most s . Thus it suffices to find the generating function $\sum_{n \geq 0} p'_s(n)x^n$, where $p'_s(n)$ denotes the number of partitions of n whose largest part is at most s . Now expanding each factor $(1 - x^k)^{-1}$ as a geometric series gives

$$\prod_{k=1}^s \frac{1}{1 - x^k} = \prod_{k=1}^s \left(\sum_{m_k \geq 0} x^{m_k k} \right).$$

How do we get a coefficient of x^n ? We must choose a term $x^{m_k k}$ from each factor of the product, $1 \leq k \leq s$, so that

$$n = \sum_{k=1}^s m_k k.$$

But such a choice is the same as choosing the partition λ of n such that the part k occurs m_k times. For instance, if $s = 4$ and we choose $m_1 = 5$, $m_2 = 0$, $m_3 = 1$, $m_4 = 2$, then we have chosen the partition $\lambda = (4, 4, 3, 1, 1, 1, 1, 1)$ of 16. Hence the coefficient of x^n is the number of partitions λ of n whose largest part is at most s , as was to be proved. \square

Note that Proposition 8.14 is “trivial” in the sense that it can be seen by inspection. There is an obvious correspondence between (a) the choice of terms contributing to the coefficient of x^n and (b) partitions of n with largest part at most r . Although the generating function we will obtain for plane partitions is equally simple, it will be far less obvious why it is correct.

Plane partitions have a certain similarity with standard Young tableaux, so perhaps it is not surprising that a variant of RSK will be applicable. Instead of NYT we will be dealing with *column-strict plane partitions* (CSPP). These are plane partitions for which the nonzero elements *strictly* decrease in each column. An example of a CSPP is given by

$$\begin{array}{cccccccc} 7 & 7 & 4 & 3 & 3 & 3 & 1 & \\ 4 & 3 & 3 & 1 & & & & \\ 3 & 2 & & & & & & \\ 2 & 1 & & & & & & \\ 1 & & & & & & & \end{array} \quad (8.17)$$

We say that this CSPP has *shape* $\lambda = (7, 4, 2, 2, 1)$, the shape of the Young diagram which the numbers occupy, and that it has five rows, seven columns, and 16 parts (so $\lambda \vdash 16$).

If $P = (P_{ij})$ is a CSPP and $k \geq 1$, then we define the *row insertion* $P \leftarrow k$ as follows: let r be the least integer such that $P_{1,r} < k$. If no such r exists (i.e., all elements of the first row of P are greater than or equal to k), then simply place k at the end of the first row. The insertion process stops, and the resulting CSPP is $P \leftarrow k$. If, on the other hand, r does exist, then replace P_{1r} by k . The element k then “bumps” $P_{1r} := k'$ into the second row, i.e., insert k' into the second row of P by the insertion rule just described, possibly bumping a new element k'' into the third row. Continue until an element is inserted at the end of a row (possibly as the first element of a new row). The resulting array is $P \leftarrow k$. Note that this rule is completely analogous to row insertion for NYT: for NYT an element bumps the leftmost element greater than it, while for CSPP an element bumps the leftmost element smaller than it.

8.15 Example. Let P be the CSPP of equation (8.17). Let us row insert 6 into P . The set of elements which get bumped are shown in bold:

$$\begin{array}{cccccc} 7 & 7 & \mathbf{4} & 3 & 3 & 3 & 1 \\ 4 & \mathbf{3} & 3 & 1 & & & \\ 3 & \mathbf{2} & & & & & \\ 2 & \mathbf{1} & & & & & \\ 1 & & & & & & \end{array} .$$

The final 1 that was bumped is inserted at the end of the fifth row. Thus we obtain

$$(P \leftarrow 6) = \begin{array}{cccccc} 7 & 7 & 6 & 3 & 3 & 3 & 1 \\ 4 & 4 & 3 & 1 & & & \\ 3 & 3 & & & & & \\ 2 & 2 & & & & & \\ 1 & 1 & & & & & \end{array} .$$

We are now ready to describe the analogue of RSK needed to count plane partitions. Instead of beginning with a permutation $\pi \in \mathfrak{S}_n$, we begin with an $r \times s$ matrix $A = (a_{ij})$ of nonnegative integers, called for short an $r \times s$ \mathbb{N} -matrix. We convert A into a *two-line array*

$$w_A = \begin{pmatrix} u_1 & u_2 & \cdots & u_N \\ v_1 & v_2 & \cdots & v_N \end{pmatrix},$$

where

- $u_1 \geq u_2 \geq \cdots \geq u_N$
- If $i < j$ and $u_i = u_j$, then $v_i \geq v_j$.
- The number of columns of w_A equal to $\begin{smallmatrix} i \\ j \end{smallmatrix}$ is a_{ij} . (It follows that $N = \sum a_{ij}$.)

It is easy to see that w_A is uniquely determined by A , and conversely. As an example, suppose that

$$A = \begin{bmatrix} 0 & 1 & 0 & 2 \\ 1 & 1 & 1 & 0 \\ 2 & 1 & 0 & 0 \end{bmatrix}. \quad (8.18)$$

Then

$$w_A = \begin{pmatrix} 3 & 3 & 3 & 2 & 2 & 2 & 1 & 1 & 1 \\ 2 & 1 & 1 & 3 & 2 & 1 & 4 & 4 & 2 \end{pmatrix}.$$

We now insert the numbers v_1, v_2, \dots, v_N successively into a CSPP. That is, we start with $P_0 = \emptyset$ and define inductively $P_i = P_{i-1} \leftarrow v_i$. We also start with $Q_0 = \emptyset$, and at the i th step insert u_i into Q_{i-1} (without any bumping or other altering of the elements of Q_{i-1}) so that P_i and Q_i have the same shape. Finally let $(P, Q) = (P_N, Q_N)$ and write $A \xrightarrow{\text{RSK}'} (P, Q)$.

8.16 Example. Let A be given by equation (8.18). The pairs $(P_1, Q_1), \dots,$


$(P_9, Q_9) = (P, Q)$ are as follows:

<u>P_i</u>	<u>Q_i</u>
2	3
2 1	3 3
2 1 1	3 3 3
3 1 1 2	3 3 3 2
3 2 1 2 1	3 3 3 2 2
3 2 1 1 2 1	3 3 3 2 2 2
4 2 1 1 3 1 2	3 3 3 2 2 2 1
4 4 1 1 3 2 2 1	3 3 3 2 2 2 1 1
4 4 2 1 3 2 1 2 1	3 3 3 2 2 2 1 1 1


It is straightforward to show that if $A \xrightarrow{\text{RSK}'} (P, Q)$, then P and Q are CSPP of the same shape. We omit the proof of the following key lemma, which is analogous to the proof of Theorem 8.13. Let us just note a crucial property (which is easy to prove) of the correspondence $A \xrightarrow{\text{RSK}'} (P, Q)$ which allows us to recover A from (P, Q) , namely, equal entries of Q are inserted from left-to-right. Thus the last number placed into Q is the rightmost occurrence of the least entry. Hence we can inverse bump the number in this position in P to back up one step in the algorithm, just as for the usual RSK correspondence $\pi \xrightarrow{\text{RSK}} (P, Q)$.

8.17 Lemma. *The correspondence $A \xrightarrow{\text{RSK}'} (P, Q)$ is a bijection from the set of $r \times s$ matrices of nonnegative integers to the set of pairs (P, Q) of CSPP of the same shape, such that the largest part of P is at most s and the largest part of Q is at most r .*


The next step is to convert the pair (P, Q) of CSPP of the same shape into a single plane partition π . We do this by “merging” the i th column of P with the i th column of Q , producing the i th column of π . Thus we first describe how to merge two partitions λ and μ with distinct parts and with the same number of parts into a single partition $\rho = \rho(\lambda, \mu)$. Draw the Ferrers diagram of λ but with each row indented one space to the right of the beginning of the previous row. Such a diagram is called the *shifted* Ferrers diagram of λ . For instance, if $\lambda = (5, 3, 2)$ then we get the shifted diagram



Do the same for μ , and then transpose the diagram. For instance, if $\mu = (6, 3, 1)$ then we get the transposed shifted diagram



Now merge the two diagrams into a single diagram by identifying their main diagonals. For λ and μ as above, we get the diagram (with the main diagonal drawn for clarity):



Define $\rho(\lambda, \mu)$ to be the partition for which this merged diagram is the Ferrers diagram. The above example shows that

$$\rho(532, 631) = 544211.$$

The map $(\lambda, \mu) \mapsto \rho(\lambda, \mu)$ is clearly a bijection between pairs of partitions (λ, μ) with k distinct parts, and partitions ρ whose main diagonal (of the Ferrers diagram) has k dots. Equivalently, k is the largest integer j for which $\rho_j \geq j$. Note that

$$|\rho| = |\lambda| + |\mu| - \ell(\lambda). \quad (8.19)$$

We now extend the above bijection to pairs (P, Q) of reverse SSYT of the same shape. If λ^i denotes the i th column of P and μ^i the i th column of Q , then let $\pi(P, Q)$ be the array whose i th column is $\rho(\lambda^i, \mu^i)$. For instance, if

$$P = \begin{array}{cccc} 4 & 4 & 2 & 1 \\ 3 & 1 & 1 & \\ 2 & & & \end{array} \quad \text{and} \quad Q = \begin{array}{cccc} 5 & 3 & 2 & 2 \\ 4 & 2 & 1 & \\ 1 & & & \end{array},$$

then

$$\pi(P, Q) = \begin{array}{cccc} 4 & 4 & 2 & 1 \\ 4 & 2 & 2 & 1 \\ 4 & 2 & & \\ 2 & & & \\ 2 & & & \end{array}.$$

It is easy to see that $\pi(P, Q)$ is a plane partition. Replace each row of $\pi(P, Q)$ by its conjugate to obtain another plane partition $\pi'(P, Q)$. With $\pi(P, Q)$ as above we obtain

$$\pi'(P, Q) = \begin{array}{cccc} 4 & 3 & 2 & 2 \\ 4 & 3 & 1 & 1 \\ 2 & 2 & 1 & 1 \\ 1 & 1 & & \\ 1 & 1 & & \end{array}.$$

Write $|P|$ for the sum of the elements of P , and write $\max(P)$ for the largest element of P , and similarly for Q . When we merge P and Q into $\pi(P, Q)$, $\max(P)$ becomes the largest part of $\pi(P, Q)$. Thus when we conjugate each row, $\max(P)$ becomes the number $\text{col}(\pi'(P, Q))$ of columns of $\pi'(P, Q)$ [why?]. Similarly, $\max(Q)$ becomes the number $\text{row}(\pi'(P, Q))$ of rows of $\pi(P, Q)$ and of $\pi'(P, Q)$. In symbols,

$$\begin{aligned} \max P &= \text{col}(\pi'(P, Q)) \\ \max Q &= \text{row}(\pi'(P, Q)). \end{aligned} \quad (8.20)$$

Moreover, it follows from equation (8.19) that

$$|\pi'(P, Q)| = |\pi(P, Q)| = |P| + |Q| - \nu(P), \quad (8.21)$$

where $\nu(P)$ denotes the number of parts of P (or of Q).

We now have all the ingredients necessary to prove the main result of this appendix.

8.18 Theorem. *Let $\text{pp}_{rs}(n)$ denote the number of plane partitions of n with at most r rows and at most s columns. Then*

$$\sum_{n \geq 0} \text{pp}_{rs}(n) x^n = \prod_{i=1}^r \prod_{j=1}^s (1 - x^{i+j-1})^{-1}.$$

Proof. Let $A = (a_{ij})$ be an $r \times s$ \mathbb{N} -matrix. We can combine the bijections discussed above to obtain a plane partition $\pi(A)$ associated with A . Namely, first apply RSK to obtain $A \xrightarrow{\text{RSK}'} (P, Q)$, and then apply the merging process and row conjugation to obtain $\pi(A) = \pi'(P, Q)$. Since a column $\begin{smallmatrix} i \\ j \end{smallmatrix}$ of the two-line array w_A occurs a_{ij} times and results in an insertion of j into P and i into Q , it follows that

$$\begin{aligned} |P| &= \sum_{i,j} j a_{ij} \\ |Q| &= \sum_{i,j} i a_{ij} \\ \max(P) &= \max\{j : a_{ij} \neq 0\} \\ \max(Q) &= \max\{i : a_{ij} \neq 0\} \end{aligned}$$

Hence from equations (8.20) and (8.21), we see that the map $A \mapsto \pi(A)$ is a bijection from $r \times s$ \mathbb{N} -matrices A to plane partitions with at most r rows and at most s columns. Moreover,

$$\begin{aligned} |\pi(A)| &= |P| + |Q| - \nu(P) \\ &= \sum_{i,j} (i + j - 1) a_{ij}. \end{aligned}$$

Thus the enumeration of plane partitions is reduced to the much easier enu-

meration of \mathbb{N} -matrices. Specifically, we have

$$\begin{aligned}
 \sum_{n \geq 0} \text{pp}_{rs}(n)x^n &= \sum_{\substack{\pi \\ \text{row}(\pi) \leq r \\ \text{col}(\pi) \leq s}} x^{|\pi|} \\
 &= \sum_{r \times s \text{ } \mathbb{N}\text{-matrices } A} x^{\sum (i+j-1)a_{ij}} \\
 &= \prod_{i=1}^r \prod_{j=1}^s \left(\sum_{a_{ij} \geq 0} x^{\sum (i+j-1)a_{ij}} \right) \\
 &= \prod_{i=1}^r \prod_{j=1}^s (1 - x^{i+j-1})^{-1}.
 \end{aligned}$$

□

Write $\text{pp}_r(n)$ for the number of plane partitions of n with at most r rows. Letting $s \rightarrow \infty$ and then $r \rightarrow \infty$ in Theorem 8.18 produces the elegant generating functions of the next corollary.

8.19 Corollary. *We have*

$$\sum_{n \geq 0} \text{pp}_r(n)x^n = \prod_{i \geq 1} (1 - x^i)^{-\min(i,r)} \quad (8.22)$$

$$\sum_{n \geq 0} \text{pp}(n)x^n = \prod_{i \geq 1} (1 - x^i)^{-i}. \quad (8.23)$$

NOTE. Once one has seen the generating function

$$\frac{1}{(1-x)(1-x^2)(1-x^3)\cdots}$$

for one-dimensional (ordinary) partitions and the generating function

$$\frac{1}{(1-x)(1-x^2)^2(1-x^3)^3\cdots}$$

for two-dimensional (plane) partitions, it is quite natural to ask about higher-dimensional partitions. In particular, a *solid partition* of n is a three-dimensional array $\pi = (\pi_{ijk})_{i,j,k \geq 1}$ of nonnegative integers, weakly decreasing in each of

the three coordinate directions, and with elements summing to n . Let $\text{sol}(n)$ denote the number of solid partitions of n . It is easy to see that for any integer sequence $a_0 = 1, a_1, a_2, \dots$, there are unique integers b_1, b_2, \dots for which

$$\sum_{n \geq 0} a_n x^n = \prod_{i \geq 1} (1 - x^i)^{-b_i}.$$

For the case $a_n = \text{sol}(n)$, we have

$$b_1 = 1, b_2 = 3, b_3 = 6, b_4 = 10, b_5 = 15,$$

which looks quite promising. Alas, the sequence of exponents continues

$$20, 26, 34, 46, 68, 97, 120, 112, 23, -186, -496, -735, -531, 779, \dots$$

The problem of enumerating solid partitions remains open and is considered most likely to be hopeless.

Notes for Chapter 8

Standard Young tableaux (SYT) were first enumerated by P. A. MacMahon [76, p. 175] (see also [77, §103]). MacMahon formulated his result in terms of “generalized ballot sequences” or “lattice permutations” rather than SYT, but they are easily seen to be equivalent. He stated the result not in terms of the products of hook lengths as in Theorem 8.1, but as a more complicated product formula. The formulation in terms of hook lengths is due to J. S. Frame and appears first in the paper [39, Thm. 1] of Frame, Robinson, and R. M. Thrall; hence it is sometimes called the “Frame-Robinson-Thrall hook-length formula.” (The actual definition of standard Young tableaux is due to A. Young [122, p. 258].)

Independently of MacMahon, F. G. Frobenius [42, eqn. (6)] obtained the same formula for the degree of the irreducible character χ^λ of \mathfrak{S}_n as MacMahon obtained for the number of lattice permutations of type λ . Frobenius was apparently unaware of the combinatorial significance of $\deg \chi^\lambda$, but Young showed in [122, pp. 260–261] that $\deg \chi^\lambda$ was the number of SYT of shape λ , thereby giving an independent proof of MacMahon’s result. (Young also provided his own proof of MacMahon’s result in [122, Thm. II].)

A number of other proofs of the hook-length formula were subsequently found. C. Greene, A. Nijenhuis, and H. S. Wilf [49] gave an elegant probabilistic proof. A proof of A. Hillman and R. Grassl [58] shows very clearly the

role of hook lengths, though the proof is not completely bijective. A bijective version was later given by C. F. Krattenthaler [65]. Completely bijective proofs of the hook-length formula were first given by D. S. Franzblau and D. Zeilberger [40] and by J. B. Remmel [95]. An exceptionally elegant bijective proof was later found by J.-C. Novelli, I. M. Pak, and A. V. Stoyanovskii [82].

The use of the operators U and D to count walks in the Hasse diagram of Young's lattice was developed independently, in a more general context, by S. Fomin [37][38] and R. Stanley [106][108]. See also [109, §3.21] for a short exposition.

The RSK algorithm (known by a variety of other names, either “correspondence” or “algorithm” in connection with some subset of the names Robinson, Schensted, and Knuth) was first described, in a rather vague form, by G. de B. Robinson [96, §5], as a tool in an attempted proof of a result now known as the “Littlewood–Richardson Rule.” The RSK algorithm was later rediscovered by C. E. Schensted (see below), but no one actually analyzed Robinson's work until this was done by M. van Leeuwen [69, §7]. It is interesting to note that Robinson says in a footnote on page 754 that “I am indebted for this association I to Mr. D. E. Littlewood.” Van Leeuwen's analysis makes it clear that “association I” gives the recording tableau Q of the RSK algorithm $\pi \xrightarrow{\text{RSK}} (P, Q)$. Thus it might be correct to say that if $\pi \in \mathfrak{S}_n$ and $\pi \xrightarrow{\text{RSK}} (P, Q)$, then the definition of P is due to Robinson, while the definition of Q is due to Littlewood.

No further work related to Robinson's construction was done until Schensted published his seminal paper [99] in 1961. (For some information about the unusual life of Schensted, see [6].) Schensted's purpose was the enumeration of permutations in \mathfrak{S}_n according to the length of their longest increasing and decreasing subsequences. According to Knuth [66, p. 726], the connection between the work of Robinson and that of Schensted was first pointed out by M.-P. Schützenberger, though as mentioned above the first person to describe this connection precisely was van Leeuwen.

Plane partitions were discovered by MacMahon in a series of papers which were not appreciated until much later. (See MacMahon's book [77, Sections IX and X] for an exposition of his results.) MacMahon's first paper dealing with plane partitions was [75]. In Article 43 of this paper he gives the definition of a plane partition (though not yet with that name). In Article 51 he conjectures that the generating function for plane partitions is the product

$$(1 - x)^{-1} (1 - x^2)^{-2} (1 - x^3)^{-3} (1 - x^4)^{-4} \dots$$

(our equation (8.23)). In Article 52 he conjectures our equation (8.22) and Theorem 8.18, finally culminating in a conjectured generating function for plane partitions of n with at most r rows, at most s columns, and with largest part at most t . (See Exercise 8.34.) MacMahon goes on in Articles 56–62 to prove his conjecture in the case of plane partitions with at most 2 rows and s columns (the case $r = 2$ of our Theorem 8.18), mentioning on page 662 that an independent solution was obtained by A. R. Forsyth. (Though a publication reference is given to Forsyth’s paper, apparently it never actually appeared.)

We will not attempt to describe MacMahon’s subsequent work on plane partitions, except to say that the culmination of his work appears in [77, Art. 495], in which he proves his main conjecture from his first paper [75] on plane partitions, viz., our Exercise 8.34. MacMahon’s proof is quite lengthy and indirect.

In 1972 E. A. Bender and D. E. Knuth [7] showed the connection between the theory of symmetric functions and the enumeration of plane partitions. They gave simple proofs based on the RSK algorithm of many results involving plane partitions, including the first bijective proof (the same proof that we give) of our Theorem 8.18.

For further aspects of Young tableaux and the related topics of symmetric functions, representation theory of the symmetric group, Grassmann varieties, etc., see the expositions of W. E. Fulton [43], B. E. Sagan [98], and R. Stanley [110, Ch. 7].

Chapter 9

The Matrix-Tree Theorem

The Matrix-Tree Theorem is a formula for the number of spanning trees of a graph in terms of the determinant of a certain matrix. We begin with the necessary graph-theoretical background. Let G be a finite graph, allowing multiple edges but not loops. (Loops could be allowed, but they turn out to be completely irrelevant.) We say that G is *connected* if there exists a walk between any two vertices of G . A *cycle* is a closed walk with no repeated vertices or edges, except for the the first and last vertex. A *tree* is a connected graph with no cycles. In particular, a tree cannot have multiple edges, since a double edge is equivalent to a cycle of length two. The three nonisomorphic trees with five vertices are shown in Figure 9.1.

A basic theorem of graph theory (whose easy proof we leave as an exercise) is the following.

9.1 Proposition. *Let G be a graph with p vertices. The following conditions are equivalent.*

- (a) G is a tree.
- (b) G is connected and has $p - 1$ edges.
- (c) G has no cycles and has $p - 1$ edges.





Figure 9.1: The three trees with five vertices

(d) There is a unique path (= walk with no repeated vertices) between any two vertices.

A *spanning subgraph* of a graph G is a graph H with the same vertex set as G , and such that every edge of H is an edge of G . If G has q edges, then the number of spanning subgraphs of G is equal to 2^q , since we can choose any subset of the edges of G to be the set of edges of H . (Note that multiple edges between the same two vertices are regarded as *distinguishable*, in accordance with the definition of a graph in Chapter 1.) A spanning subgraph which is a tree is called a *spanning tree*. Clearly G has a spanning tree if and only if it is connected [why?]. An important invariant of a graph G is its number of spanning trees, called the *complexity* of G and denoted $\kappa(G)$.

9.2 Example. Let G be the graph illustrated below, with edges a, b, c, d, e .



Then G has eight spanning trees, namely, $abc, abd, acd, bcd, abe, ace, bde,$ and cde (where, e.g., abc denotes the spanning subgraph with edge set $\{a, b, c\}$).

9.3 Example. Let $G = K_5$, the complete graph on five vertices. A simple counting argument shows that K_5 has 60 spanning trees isomorphic to the first tree in Figure 9.1, 60 isomorphic to the second tree, and 5 isomorphic to the third tree. Hence $\kappa(K_5) = 125$. It is even easier to verify that $\kappa(K_1) = 1$, $\kappa(K_2) = 1$, $\kappa(K_3) = 3$, and $\kappa(K_4) = 16$. Can the reader make a conjecture about the value of $\kappa(K_p)$ for any $p \geq 1$?

Our object is to obtain a “determinantal formula” for $\kappa(G)$. For this we need an important result from matrix theory, known as the *Binet-Cauchy theorem* or *Cauchy-Binet theorem* and which is often omitted from a beginning linear algebra course. Later (Theorem 10.4) we will prove a more

general determinantal formula without the use of the Binet-Cauchy theorem. However, the use of the Binet-Cauchy theorem does afford some additional algebraic insight. The Binet-Cauchy theorem is a generalization of the familiar fact that if A and B are $n \times n$ matrices, then $\det AB = (\det A)(\det B)$, where \det denotes determinant. We want to extend this formula to the case where A and B are rectangular matrices whose product is a square matrix (so that $\det AB$ is defined). In other words, A will be an $m \times n$ matrix and B an $n \times m$ matrix, for some $m, n \geq 1$.

We will use the following notation involving submatrices. Suppose $A = (a_{ij})$ is an $m \times n$ matrix, with $1 \leq i \leq m$, $1 \leq j \leq n$, and $m \leq n$. Given an m -element subset S of $\{1, 2, \dots, n\}$, let $A[S]$ denote the $m \times m$ submatrix of A obtained by taking the columns indexed by the elements of S . In other words, if the elements of S are given by $j_1 < j_2 < \dots < j_m$, then $A[S] = (a_{i,j_k})$, where $1 \leq i \leq m$ and $1 \leq k \leq m$. For instance, if

$$A = \begin{bmatrix} 1 & 2 & 3 & 4 & 5 \\ 6 & 7 & 8 & 9 & 10 \\ 11 & 12 & 13 & 14 & 15 \end{bmatrix}$$

and $S = \{2, 3, 5\}$, then

$$A[S] = \begin{bmatrix} 2 & 3 & 5 \\ 7 & 8 & 10 \\ 12 & 13 & 15 \end{bmatrix}.$$

Similarly, let $B = (b_{ij})$ be an $n \times m$ matrix with $1 \leq i \leq n$, $1 \leq j \leq m$ and $m \leq n$. Let S be an m -element subset of $\{1, 2, \dots, n\}$ as above. Then $B[S]$ denotes the $m \times m$ matrix obtained by taking the rows of B indexed by S . Note that $A^t[S] = A[S]^t$, where t denotes transpose.

9.4 Theorem (the Binet-Cauchy Theorem). *Let $A = (a_{ij})$ be an $m \times n$ matrix, with $1 \leq i \leq m$ and $1 \leq j \leq n$. Let $B = (b_{ij})$ be an $n \times m$ matrix with $1 \leq i \leq n$ and $1 \leq j \leq m$. (Thus AB is an $m \times m$ matrix.) If $m > n$, then $\det(AB) = 0$. If $m \leq n$, then*

$$\det(AB) = \sum_S (\det A[S])(\det B[S]),$$

where S ranges over all m -element subsets of $\{1, 2, \dots, n\}$.

Before proceeding to the proof, let us give an example. We write $|a_{ij}|$ for the determinant of the matrix (a_{ij}) . Suppose

$$A = \begin{bmatrix} a_1 & a_2 & a_3 \\ b_1 & b_2 & b_3 \end{bmatrix}, \quad B = \begin{bmatrix} c_1 & d_1 \\ c_2 & d_2 \\ c_3 & d_3 \end{bmatrix}.$$

Then

$$\det AB = \begin{vmatrix} a_1 & a_2 \\ b_1 & b_2 \end{vmatrix} \cdot \begin{vmatrix} c_1 & d_1 \\ c_2 & d_2 \end{vmatrix} + \begin{vmatrix} a_1 & a_3 \\ b_1 & b_3 \end{vmatrix} \cdot \begin{vmatrix} c_1 & d_1 \\ c_3 & d_3 \end{vmatrix} + \begin{vmatrix} a_2 & a_3 \\ b_2 & b_3 \end{vmatrix} \cdot \begin{vmatrix} c_2 & d_2 \\ c_3 & d_3 \end{vmatrix}.$$

Proof of Theorem 9.4 (sketch). First suppose $m > n$. Since from linear algebra we know that $\text{rank } AB \leq \text{rank } A$ and that the rank of an $m \times n$ matrix cannot exceed n (or m), we have that $\text{rank } AB \leq n < m$. But AB is an $m \times m$ matrix, so $\det AB = 0$, as claimed.

Now assume $m \leq n$. We use notation such as M_{rs} to denote an $r \times s$ matrix M . It is an immediate consequence of the definition of matrix multiplication (which the reader should check) that

$$\begin{bmatrix} R_{mm} & S_{mn} \\ T_{nm} & U_{nn} \end{bmatrix} \begin{bmatrix} V_{mn} & W_{mm} \\ X_{nn} & Y_{nm} \end{bmatrix} = \begin{bmatrix} RV + SX & RW + SY \\ TV + UX & TW + UY \end{bmatrix}. \quad (9.1)$$

In other words, we can multiply “block” matrices of suitable dimensions as if their entries were numbers. Note that the entries of the right-hand side of (9.1) all have well-defined dimensions (sizes), e.g., $RV + SX$ is an $m \times n$ matrix since both RV and SX are $m \times n$ matrices.

Now in equation (9.1) let $R = I_m$ (the $m \times m$ identity matrix), $S = A$, $T = O_{nm}$ (the $n \times m$ matrix of 0’s), $U = I_n$, $V = A$, $W = O_{mm}$, $X = -I_n$, and $Y = B$. We get

$$\begin{bmatrix} I_m & A \\ O_{nm} & I_n \end{bmatrix} \begin{bmatrix} A & O_{mm} \\ -I_n & B \end{bmatrix} = \begin{bmatrix} O_{mn} & AB \\ -I_n & B \end{bmatrix}. \quad (9.2)$$

Take the determinant of both sides of (9.2). The first matrix on the left-hand side is upper triangular with 1’s on the main diagonal. Hence its determinant is one. Since the determinant of a product of square matrices is the product of the determinants of the factors, we get

$$\begin{vmatrix} A & O_{mm} \\ -I_n & B \end{vmatrix} = \begin{vmatrix} O_{mn} & AB \\ -I_n & B \end{vmatrix}. \quad (9.3)$$

It is easy to see [why?] that the determinant on the right-hand side of (9.3) is equal to $\pm \det AB$. So consider the left-hand side. A nonzero term in the expansion of the determinant on the left-hand side is obtained by taking the product (with a certain sign) of $m + n$ nonzero entries, no two in the same row and column (so one in each row and each column). In particular, we must choose m entries from the last m columns. These entries belong to m of the bottom n rows [why?], say rows $m + s_1, m + s_2, \dots, m + s_m$. Let $S = \{s_1, s_2, \dots, s_m\} \subseteq \{1, 2, \dots, n\}$. We must choose $n - m$ further entries from the last n rows, and we have no choice but to choose the -1 's in those rows $m + i$ for which $i \notin S$. Thus every term in the expansion of the left-hand side of (9.3) uses exactly $n - m$ of the -1 's in the bottom left block $-I_n$.

What is the contribution to the expansion of the left-hand side of (9.3) from those terms which use exactly the -1 's from rows $m + i$ where $i \notin S$? We obtain this contribution by deleting all rows and columns to which these -1 's belong (in other words, delete row $m + i$ and column i whenever $i \in \{1, 2, \dots, n\} - S$), taking the determinant of the $2m \times 2m$ matrix M_S that remains, and multiplying by an appropriate sign [why?]. But the matrix M_S is in block-diagonal form, with the first block just the matrix $A[S]$ and the second block just $B[S]$. Hence $\det M_S = (\det A[S])(\det B[S])$ [why?]. Taking all possible subsets S gives


$$\det AB = \sum_{\substack{S \subseteq \{1, 2, \dots, n\} \\ |S| = m}} \pm (\det A[S])(\det B[S]).$$

It is straightforward but somewhat tedious to verify that all the signs are $+$; we omit the details. This completes the proof. \square

In Chapter 1 we defined the adjacency matrix $\mathbf{A}(G)$ of a graph G with vertex set $V = \{v_1, \dots, v_p\}$ and edge set $E = \{e_1, \dots, e_q\}$. We now define two related matrices. Continue to assume that G has no loops. (This assumption is harmless since loops have no effect on $\kappa(G)$.)

9.5 Definition. Let G be as above. Give G an *orientation* \mathfrak{o} , i.e., for every edge e with vertices u, v , choose one of the ordered pairs (u, v) or (v, u) . If we choose (u, v) , say, then we think of putting an arrow on e pointing from u to v ; and we say that e is directed from u to v , that u is the *initial vertex* and v the *final vertex* of e , etc.

(a) The *incidence matrix* $\mathbf{M}(G)$ of G (with respect to the orientation \mathfrak{o})

Figure 9.2: A graph G with an orientation \mathfrak{o}

is the $p \times q$ matrix whose (i, j) -entry M_{ij} is given by

$$M_{ij} = \begin{cases} -1, & \text{if the edge } e_j \text{ has initial vertex } v_i \\ 1, & \text{if the edge } e_j \text{ has final vertex } v_i \\ 0, & \text{otherwise.} \end{cases}$$

(b) The *laplacian matrix* $L(G)$ of G is the $p \times p$ matrix whose (i, j) -entry L_{ij} is given by

$$L_{ij} = \begin{cases} -m_{ij}, & \text{if } i \neq j \text{ and there are } m_{ij} \text{ edges between } v_i \text{ and } v_j \\ \deg(v_i), & \text{if } i = j, \end{cases}$$

where $\deg(v_i)$ is the number of edges incident to v_i . Note that $L(G)$ is symmetric and does not depend on the orientation \mathfrak{o} .

As an example, let (G, \mathfrak{o}) be the oriented graph of Figure 9.2. Then

$$M(G) = \begin{bmatrix} 1 & -1 & 0 & -1 & -1 & -1 \\ -1 & 1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix}$$

$$L(G) = \begin{bmatrix} 4 & -2 & -1 & -2 \\ -2 & 3 & -1 & 0 \\ -1 & -1 & 2 & 0 \\ -2 & 0 & 0 & 2 \end{bmatrix}.$$

For any graph G , every column of $M(G)$ contains one 1, one -1 , and $q - 2$ 0's; and hence the sum of the entries in each column is 0. Thus all

the rows sum to the 0 vector, a linear dependence relation which shows that $\text{rank}(\mathbf{M}(G)) < p$. Two further properties of $\mathbf{M}(G)$ and $\mathbf{L}(G)$ are given by the following lemma.

9.6 Lemma. (a) We have $\mathbf{M}\mathbf{M}^t = \mathbf{L}$.

(b) If G is regular of degree d , then $\mathbf{L}(G) = dI - \mathbf{A}(G)$, where $\mathbf{A}(G)$ denotes the adjacency matrix of G . Hence if G (or $\mathbf{A}(G)$) has eigenvalues $\lambda_1, \dots, \lambda_p$, then $\mathbf{L}(G)$ has eigenvalues $d - \lambda_1, \dots, d - \lambda_p$.

Proof. (a) This is immediate from the definition of matrix multiplication. Specifically, for $v_i, v_j \in V(G)$ we have

$$(\mathbf{M}\mathbf{M}^t)_{ij} = \sum_{e_k \in E(G)} \mathbf{M}_{ik}\mathbf{M}_{jk}.$$

If $i \neq j$, then in order for $\mathbf{M}_{ik}\mathbf{M}_{jk} \neq 0$, we must have that the edge e_k connects the vertices v_i and v_j . If this is the case, then one of \mathbf{M}_{ik} and \mathbf{M}_{jk} will be 1 and the other -1 [why?], so their product is always -1 . Hence $(\mathbf{M}\mathbf{M}^t)_{ij} = -m_{ij}$, as claimed.

There remains the case $i = j$. Then $\mathbf{M}_{ik}\mathbf{M}_{ik}$ will be 1 if e_k is an edge with v_i as one of its vertices and will be 0 otherwise [why?]. So now we get $(\mathbf{M}\mathbf{M}^t)_{ii} = \text{deg}(v_i)$, as claimed. This proves (a).

(b) Clear by (a), since the diagonal elements of $\mathbf{M}\mathbf{M}^t$ are all equal to d . \square

Now assume that G is connected, and let $\mathbf{M}_0(G)$ be $\mathbf{M}(G)$ with its last row removed. Thus $\mathbf{M}_0(G)$ has $p - 1$ rows and q columns. Note that the number of rows is equal to the number of edges in a spanning tree of G . We call $\mathbf{M}_0(G)$ the *reduced incidence matrix* of G . The next result tells us the determinants (up to sign) of all $(p - 1) \times (p - 1)$ submatrices N of \mathbf{M}_0 . Such submatrices are obtained by choosing a set $X = \{e_{i_1}, \dots, e_{i_{p-1}}\}$ of $p - 1$ edges of G , and taking all columns of \mathbf{M}_0 indexed by the set $S = \{i_1, \dots, i_{p-1}\}$. Thus this submatrix is just $\mathbf{M}_0[S]$. For convenience we will not bother to distinguish between the set S of indices with the corresponding set X of edges.

9.7 Lemma. Let S be a set of $p - 1$ edges of G . If S does not form the set of edges of a spanning tree, then $\det \mathbf{M}_0[S] = 0$. If, on the other hand, S is the set of edges of a spanning tree of G , then $\det \mathbf{M}_0[S] = \pm 1$.

Proof. If S is not the set of edges of a spanning tree, then some subset R of S forms the edges of a cycle C in G . Suppose that the cycle C defined by R has edges f_1, \dots, f_s in that order. Multiply the column of $\mathbf{M}_0[S]$ indexed by f_j by 1 if in going around C we traverse f_i in the direction of its arrow; otherwise multiply the column by -1 . Then add these modified columns. It is easy to see (check a few small examples to convince yourself) that we get the 0 column. Hence the columns of $\mathbf{M}_0[S]$ are linearly dependent, so $\det \mathbf{M}_0[S] = 0$, as claimed.

Now suppose that S is the set of edges of a spanning tree T . Let e be an edge of T which is connected to v_p (the vertex which indexed the bottom row of \mathbf{M} , i.e., the row removed to get \mathbf{M}_0). The column of $\mathbf{M}_0[S]$ indexed by e contains exactly one nonzero entry [why?], which is ± 1 . Remove from $\mathbf{M}_0[S]$ the row and column containing the nonzero entry of column e , obtaining a $(p-2) \times (p-2)$ matrix \mathbf{M}'_0 . Note that $\det \mathbf{M}_0[S] = \pm \det \mathbf{M}'_0$ [why?]. Let T' be the tree obtained from T by contracting the edge e to a single vertex (so that v_p and the remaining vertex of e are merged into a single vertex u). Then \mathbf{M}'_0 is just the matrix obtained from the incidence matrix $\mathbf{M}(T')$ by removing the row indexed by u [why?]. Hence by induction on the number p of vertices (the case $p = 1$ being trivial), we have $\det \mathbf{M}'_0 = \pm 1$. Thus $\det \mathbf{M}_0[S] = \pm 1$, and the proof follows. \square

NOTE. An alternative way of seeing that $\det \mathbf{M}_0[S] = \pm 1$ when S is the set of edges of a spanning tree T is as follows. Let u_1, u_2, \dots, u_{p-1} be an ordering of the vertices v_1, \dots, v_{p-1} such that u_i is an endpoint of the tree obtained from T by removing vertices u_1, \dots, u_{i-1} . (It is easy to see that such an ordering is possible.) Permute the rows of $\mathbf{M}_0[S]$ so that the i th row is indexed by u_i . Then permute the columns in the order e_1, \dots, e_{p-1} so that e_i is the unique edge adjacent to u_i after u_1, \dots, u_{i-1} have been removed. Then we obtain a lower triangular matrix with ± 1 's on the main diagonal, so the determinant is ± 1 .

We have now assembled all the ingredients for the main result of this chapter. Recall that $\kappa(G)$ denotes the number of spanning trees of G .

9.8 Theorem (the Matrix-Tree Theorem). *Let G be a finite connected graph without loops, with laplacian matrix $\mathbf{L} = \mathbf{L}(G)$. Let \mathbf{L}_0 denote \mathbf{L} with the last row and column removed (or with the i th row and column removed for any i). Then*

$$\det \mathbf{L}_0 = \kappa(G).$$

Proof. Since $\mathbf{L} = \mathbf{M}\mathbf{M}^t$ (Lemma 9.6(a)), it follows immediately that $\mathbf{L}_0 = \mathbf{M}_0\mathbf{M}_0^t$. Hence by the Binet-Cauchy theorem (Theorem 9.4), we have

$$\det \mathbf{L}_0 = \sum_S (\det \mathbf{M}_0[S]) (\det \mathbf{M}_0^t[S]), \quad (9.4)$$

where S ranges over all $(p-1)$ -element subsets of $\{1, 2, \dots, q\}$ (or equivalently, over all $(p-1)$ -element subsets of the set of edges of G). Since in general $A^t[S] = A[S]^t$, equation (9.4) becomes

$$\det \mathbf{L}_0 = \sum_S (\det \mathbf{M}_0[S])^2. \quad (9.5)$$

According to Lemma 9.7, $\det \mathbf{M}_0[S]$ is ± 1 if S forms the set of edges of a spanning tree of G , and is 0 otherwise. Therefore the term indexed by S in the sum on the right-hand side of (9.5) is 1 if S forms the set of edges of a spanning tree of G , and is 0 otherwise. Hence the sum is equal to $\kappa(G)$, as desired. \square

The operation of removing a row and column from $\mathbf{L}(G)$ may seem somewhat contrived. We would prefer a description of $\kappa(G)$ directly in terms of $\mathbf{L}(G)$. Such a description will follow from the next lemma.

9.9 Lemma. *Let M be a $p \times p$ matrix (with entries in a field) such that the sum of the entries in every row and column is 0. Let M_0 be the matrix obtained from M by removing the last row and last column (or more generally, any row and any column). Then the coefficient of x in the characteristic polynomial $\det(M - xI)$ of M is equal to $-p \cdot \det(M_0)$. (Moreover, the constant term of $\det(M - xI)$ is 0.)*

Proof. The constant term of $\det(M - xI)$ is $\det M$, which is 0 since the rows of M sum to 0.

For simplicity we prove the rest of the lemma only for removing the last row and column, though the proof works just as well for any row and column. Add all the rows of $M - xI$ except the last row to the last row. This doesn't affect the determinant, and will change the entries of the last row all to $-x$ (since the rows of M sum to 0). Factor out $-x$ from the last row, yielding a matrix $N(x)$ satisfying $\det(M - xI) = -x \det N(x)$. Hence the coefficient of x in $\det(M - xI)$ is given by $-\det N(0)$. Now add all the columns of $N(0)$ except the last column to the last column. This does not affect $\det N(0)$.

Because the columns of M sum to 0, the last column of $N(0)$ becomes the column vector $[0, 0, \dots, 0, p]^t$. Expanding the determinant by the last column shows that $\det N(0) = p \cdot \det M_0$, and the proof follows. \square

9.10 Corollary. (a) Let G be a connected (loopless) graph with p vertices. Suppose that the eigenvalues of $\mathbf{L}(G)$ are $\mu_1, \dots, \mu_{p-1}, \mu_p$, with $\mu_p = 0$. Then

$$\kappa(G) = \frac{1}{p} \mu_1 \mu_2 \cdots \mu_{p-1}.$$

(b) Suppose that G is also regular of degree d , and that the eigenvalues of $\mathbf{A}(G)$ are $\lambda_1, \dots, \lambda_{p-1}, \lambda_p$, with $\lambda_p = d$. Then

$$\kappa(G) = \frac{1}{p} (d - \lambda_1)(d - \lambda_2) \cdots (d - \lambda_{p-1}).$$

Proof. (a) We have

$$\begin{aligned} \det(\mathbf{L} - xI) &= (\mu_1 - x) \cdots (\mu_{p-1} - x)(\mu_p - x) \\ &= -(\mu_1 - x)(\mu_2 - x) \cdots (\mu_{p-1} - x)x. \end{aligned}$$

Hence the coefficient of x is $-\mu_1 \mu_2 \cdots \mu_{p-1}$. By Lemma 9.9, we get $-\mu_1 \mu_2 \cdots \mu_{p-1} = p \cdot \det(\mathbf{L}_0)$. By Theorem 9.8 we have $\det(\mathbf{L}_0) = \kappa(G)$, and the proof follows.

(b) Immediate from (a) and Lemma 9.6(b). \square

Let us look at a couple of examples of the use of the Matrix-Tree Theorem.

9.11 Example. Let $G = K_p$, the complete graph on p vertices. Now K_p is regular of degree $d = p - 1$, and by Proposition 1.5 its eigenvalues are -1 ($p - 1$ times) and $p - 1 = d$. Hence from Corollary 9.10 there follows

$$\kappa(K_p) = \frac{1}{p} ((p - 1) - (-1))^{p-1} = p^{p-2}.$$

Naturally a combinatorial proof of such an elegant result is desirable. In the Appendix to this chapter we give three such proofs.

9.12 Example. Let $G = C_n$, the n -cube discussed in Chapter 2. Now C_n is regular of degree n , and by Corollary 2.4 its eigenvalues are $n - 2i$ with

multiplicity $\binom{n}{i}$ for $0 \leq i \leq n$. Hence from Corollary 9.10 there follows the amazing result


$$\begin{aligned}\kappa(C_n) &= \frac{1}{2^n} \prod_{i=1}^n (2i)^{\binom{n}{i}} \\ &= 2^{2^n - n - 1} \prod_{i=1}^n i^{\binom{n}{i}}.\end{aligned}$$

A direct combinatorial proof (though not an explicit bijection) was found by O. Bernardi in 2012.

APPENDIX: THREE ELEGANT COMBINATORIAL PROOFS

In this appendix we give three elegant combinatorial proofs that the number of spanning trees of the complete graph K_p is p^{p-2} (Example 9.11). The proofs are given in chronological order of their discovery.

First proof (Prüfer). Given a spanning tree T of K_p , i.e., a tree on the vertex set $[p]$, remove the largest endpoint (leaf) v and write down the vertex a_1 adjacent to v . Continue this procedure until only two vertices remain, obtaining a sequence $(a_1, \dots, a_{p-2}) \in [p]^{p-2}$, called the *Prüfer sequence* of T . For the tree below, we first remove 11 and then record 8. Next remove 10 and record 1. Then remove 8 and record 4, etc., ending with the sequence $(8, 1, 4, 4, 1, 4, 9, 1, 9)$ and leaving the two vertices 1 and 9.



We claim that the map just defined from trees T on $[p]$ to sequences $(a_1, \dots, a_{p-2}) \in [p]^{p-2}$ is a bijection, thereby completing the proof since clearly $[p]^{p-2}$ has p^{p-2} elements. The crucial observation is that the first vertex to be removed from T is the largest vertex of T missing from the sequence [why? — this takes a little thought]. This vertex is adjacent to a_1 . For our example, we get that 11 was the first vertex removed, and that 11 is adjacent to 8. We can now proceed recursively. If T_1 denotes T with the largest missing vertex removed, then the Prüfer sequence of T_1 is (a_2, \dots, a_{p-2}) . The first vertex to be removed from T_1 is the largest vertex of T_1 missing from (a_2, \dots, a_{p-2}) . This missing vertex is adjacent to a_2 . For our example, this missing vertex is 10 (since 11 is not a vertex of T_1), which is adjacent to 1. Continuing in this way, we determine one new edge of T at each step. At the end we have found $p - 2$ edges, and the remaining two unremoved vertices form the $(p - 1)$ st edge.


Second proof (Joyal). A *doubly-rooted tree* is a tree T with one vertex u labelled S (for “start”) and one vertex v (which may equal u) labelled E

(“end”). Let $t(p)$ be the number of trees T on the vertex set $[p]$, and let $d(p)$ be the number of doubly-rooted trees on $[p]$. Thus


$$d(p) = p^2 t(p), \tag{9.6}$$

since once we have chosen T there are p choices for u and p choices for v .


Let T be a doubly-rooted tree. There is a unique path from S to E , say with vertices $S = b_1, b_2, \dots, b_k = E$ (in that order). The following diagram shows such a doubly-rooted tree.



Let $a_1 < a_2 < \dots < a_k$ be the increasing rearrangement of the numbers b_1, b_2, \dots, b_k . Let π be the permutation of the set $\{a_1, \dots, a_k\}$ given by $\pi(a_i) = b_i$. Let D_π be the digraph of π , that is, the vertex set of D_π is $\{a_1, \dots, a_k\}$, with a directed edge $a_i \rightarrow b_i$ for $1 \leq i \leq k$. Since any permutation π of a finite set is a disjoint product of cycles, it follows that D_π is a disjoint union of directed cycles (all edges of each cycle point in the same direction as we traverse the cycle). For the example above, we have $k = 7$, $(b_1, \dots, b_7) = (11, 10, 15, 7, 5, 2, 3)$ and $(a_1, \dots, a_7) = (2, 3, 5, 7, 10, 11, 15)$. The digraph D_π is shown below.




Now attach to each vertex v of D_π the same subgraph T_v that was attached “below” v in T , and direct the edges of T_v toward v , obtaining a digraph D_T . For our example we get



The graph D_T has the crucial property that every vertex has outdegree one, that is, one arrow pointing out. In other words, D_T is the graph of a function $f: [p] \rightarrow [p]$, with vertex set $[p]$ and edges $i \rightarrow f(i)$. Conversely, given a function $f: [p] \rightarrow [p]$, all the above steps can be reversed to obtain a unique doubly-rooted tree T for which D_T is the graph of f . We have therefore found a bijection from doubly-rooted trees on $[p]$ to functions $f: [p] \rightarrow [p]$. Since the number of such functions f is p^p , it follows that $d(p) = p^p$. Then from equation (9.6) we get $t(p) = p^{p-2}$.

Third proof (Pitman). A *forest* is a graph without cycles; thus every connected component is a tree. A *planted forest* is a forest F for which every component T has a distinguished vertex r_T (called the *root* of T). Thus if a component T has k vertices, then there are k ways to choose the root of T .

Let P_p be the set of all planted forests on $[p]$. Let uv be an edge of a forest $F \in P_p$ such that u is closer than v to the root r of its component. Define F to *cover* the rooted forest F' if F' is obtained by removing the edge uv from F , and rooting the new tree containing v at v . This definition of cover defines the covering relation of a partial order on P_p . Under this partial order P_p is graded of rank $p - 1$. The rank of a forest F in P_p is its number of edges. The following diagram shows the poset P_3 , with the root of each tree being its top vertex.



It is an easy exercise to see that an element F of P_p of rank i covers i elements and is covered by $(p-i-1)p$ elements. We now count in two ways the number M_p of maximal chains of P_p . On the one hand, we can start at the top. The number of maximal elements of P_p is $p \cdot t(p)$, where $t(p)$ as above is the number of trees on the vertex set $[p]$, since there are p ways to choose the root of such a tree. Once a maximal element F is chosen, then there are $p-1$ elements F' that it covers, then $p-2$ elements that F' covers, etc., giving

$$M_p = p \cdot t(p)(p-1)! = p! t(p). \quad (9.7)$$

On the other hand, we can start at the bottom. There is a unique element F of rank one (the planted forest with no edges), then $(p-1)p$ elements F' that cover F , then $(p-2)p$ elements that cover F' , etc., giving

$$M_p = p^{p-1}(p-1)!. \quad (9.8)$$

Comparing equations (9.7) and (9.8) gives $t(p) = p^{p-2}$.

Our third proof isn't an explicit bijection like the first two proofs. On the other hand, it has the virtue of not depending on the names of the vertices. Note that in the first two proofs it is necessary to know when one vertex is larger than another.

Notes for Chapter 9

The concept of tree as a formal mathematical object goes back to G. Kirchhoff and K. G. C. von Staudt. Trees were first extensively investigated by A. Cayley, to whom the term “tree” is due. In particular, in [24] Cayley states the formula $\kappa(K_p) = p^{p-2}$ for the number of spanning trees of K_p , and he gives a vague idea of a combinatorial proof. Because of this paper, Cayley is often credited with the enumeration of labelled trees. Cayley pointed out, however, that an equivalent result had been proved earlier by C. W. Borchardt [11]. Moreover, this result appeared even earlier in a paper of J. J. Sylvester [113]. Undoubtedly Cayley and Sylvester could have furnished a complete, rigorous proof had they the inclination to do so. The elegant combinatorial proofs given in the appendix are due to E. P. H. Prüfer [91], A. Joyal [62, Exam. 12, pp. 15–16] and J. W. Pitman [85].

The Matrix-Tree Theorem (Theorem 9.8) was first proved by C. W. Borchardt [11] in 1860, though a similar result had earlier been published by J. J. Sylvester [113] in 1857. Cayley [23, p. 279] in fact in 1856 referred to the not-yet-published work of Sylvester. For further historical information on the Matrix-Tree theorem, see Moon [79, p. 42].

Chapter 10

Eulerian digraphs and oriented trees

A famous problem which goes back to Euler asks for what graphs G is there a closed walk which uses every edge exactly once. (There is also a version for non-closed walks.) Such a walk is called an *Eulerian tour* (also known as an *Eulerian cycle*). A graph which has an Eulerian tour is called an *Eulerian graph*. Euler's famous theorem (the first real theorem of graph theory) states that a graph G without isolated vertices (which clearly would be irrelevant) is Eulerian if and only if it is connected and every vertex has even degree. Here we will be concerned with the analogous theorem for directed graphs. We want to know not just whether an Eulerian tour exists, but how many there are. We will prove an elegant determinantal formula for this number closely related to the Matrix-Tree Theorem. For the case of undirected graphs no analogous formula is known, explaining why we consider only the directed case.

A (finite) *directed graph* or *digraph* D consists of a vertex set $V = \{v_1, \dots, v_p\}$ and edge set $E = \{e_1, \dots, e_q\}$, together with a function $\varphi: E \rightarrow V \times V$ (the set of ordered pairs (u, v) of elements of V). If $\varphi(e) = (u, v)$, then we think of e as an arrow from u to v . We then call u the *initial vertex* and v the *final vertex* of e . (These concepts arose in the definition of an orientation in Definition 8.5.) A *tour* in D is a sequence e_1, e_2, \dots, e_r of *distinct* edges such that the final vertex of e_i is the initial vertex of e_{i+1} for all $1 \leq i \leq r - 1$, and the final vertex of e_r is the initial vertex of e_1 . A tour is *Eulerian* if every edge of D occurs at least once (and hence exactly once). A digraph which has no isolated vertices and contains an Eulerian tour is

called an *Eulerian digraph*. Clearly an Eulerian digraph is connected. The *outdegree* of a vertex v , denoted $\text{outdeg}(v)$, is the number of edges of D with initial vertex v . Similarly the *indegree* of v , denoted $\text{indeg}(v)$, is the number of edges of D with final vertex v . A loop (edge e for which $\varphi(e) = (v, v)$) contributes one to both the indegree and outdegree. A digraph is *balanced* if $\text{indeg}(v) = \text{outdeg}(v)$ for all vertices v .

10.1 Theorem. *A digraph D without isolated vertices is Eulerian if and only if it is connected and balanced.*

Proof. Assume D is Eulerian, and let e_1, \dots, e_q be an Eulerian tour. As we move along the tour, whenever we enter a vertex v we must exit it, except at the very end we enter the final vertex v of e_q without exiting it. However, at the beginning we exited v without having entered it. Hence every vertex is entered as often as it is exited and so must have the same outdegree as indegree. Therefore D is balanced, and as noted above D is clearly connected.

Now assume that D is balanced and connected. We may assume that D has at least one edge. We first claim that for any edge e of D , D has a tour for which $e = e_1$. If e_1 is a loop we are done. Otherwise we have entered the vertex $\text{fin}(e_1)$ for the first time, so since D is balanced there is some exit edge e_2 . Either $\text{fin}(e_2) = \text{init}(e_1)$ and we are done, or else we have entered the vertex $\text{fin}(e_2)$ once more than we have exited it. Since D is balanced there is new edge e_3 with $\text{fin}(e_2) = \text{init}(e_3)$. Continuing in this way, either we complete a tour or else we have entered the current vertex once more than we have exited it, in which case we can exit along a new edge. Since D has finitely many edges, eventually we must complete a tour. Thus D does have a tour which uses e_1 .

Now let e_1, \dots, e_r be a tour C of maximum length. We must show that $r = q$, the number of edges of D . Assume to the contrary that $r < q$. Since in moving along C every vertex is entered as often as it is exited (with $\text{init}(e_1)$ exited at the beginning and entered at the end), when we remove the edges of C from D we obtain a digraph H which is still balanced, though it need not be connected. However, since D is connected, at least one connected component H_1 of H contains at least one edge and has a vertex v in common with C [why?]. Since H_1 is balanced, there is an edge e of H_1 with initial vertex v . The argument of the previous paragraph shows that H_1 has a tour C' of positive length beginning with the edge e . But then when moving along C , when we reach v we can take the “detour” C' before continuing with C .

This gives a tour of length longer than r , a contradiction. Hence $r = q$, and the theorem is proved. \square

Our primary goal is to count the number of Eulerian tours of a connected balanced digraph. A key concept in doing so is that of an oriented tree. An *oriented tree* with root v is a (finite) digraph T with v as one of its vertices, such that there is a unique directed path from any vertex u to v . In other words, there is a unique sequence of edges e_1, \dots, e_r such that (a) $\text{init}(e_1) = u$, (b) $\text{fin}(e_r) = v$, and (c) $\text{fin}(e_i) = \text{init}(e_{i+1})$ for $1 \leq i \leq r - 1$. It's easy to see that this means that the underlying undirected graph (i.e., “erase” all the arrows from the edges of T) is a tree, and that all arrows in T “point toward” v . There is a surprising connection between Eulerian tours and oriented trees, given by the next result.

10.2 Theorem. *Let D be a connected balanced digraph with vertex set V . Fix an edge e of D , and let $v = \text{init}(e)$. Let $\tau(D, v)$ denote the number of oriented (spanning) subtrees of D with root v , and let $\epsilon(D, e)$ denote the number of Eulerian tours of D starting with the edge e . Then*

$$\epsilon(D, e) = \tau(D, v) \prod_{u \in V} (\text{outdeg}(u) - 1)!. \quad (10.1)$$

Proof. Let $e = e_1, e_2, \dots, e_q$ be an Eulerian tour \mathcal{E} in D . For each vertex $u \neq v$, let $e(u)$ be the “last exit” from u in the tour, i.e., let $e(u) = e_j$ where $\text{init}(e(u)) = u$ and $\text{init}(e_k) \neq u$ for any $k > j$.

Claim #1. The vertices of D , together with the edges $e(u)$ for all vertices $u \neq v$, form an oriented subtree of D with root v .

Proof of Claim #1. This is a straightforward verification. Let T be the spanning subgraph of D with edges $e(u)$, $u \neq v$. Thus if $\#V = p$, then T has p vertices and $p - 1$ edges [why?]. There are three items to check to insure that T is an oriented tree with root v :

- (a) T does not have two edges f and f' satisfying $\text{init}(f) = \text{init}(f')$. This is clear since both f and f' can't be last exits from the same vertex.
- (b) T does not have an edge f with $\text{init}(f) = v$. This is clear since by definition the edges of T consist only of last exits from vertices other than v , so no edge of T can exit from v .

- (c) T does not have a (directed) cycle C . For suppose C were such a cycle. Let f be that edge of C which occurs after all the other edges of C in the Eulerian tour \mathcal{E} . Let f' be the edge of C satisfying $\text{fin}(f) = \text{init}(f')$ ($= u$, say). We can't have $u = v$ by (b). Thus when we enter u via f , we must exit u . We can't exit u via f' since f occurs after f' in \mathcal{E} . Hence f' is not the last exit from u , contradicting the definition of T .

It's easy to see that conditions (a)–(c) imply that T is an oriented tree with root v , proving the claim.

Claim #2. We claim that the following converse to Claim #1 is true. Given a connected balanced digraph D and a vertex v , let T be an oriented (spanning) subtree of D with root v . Then we can construct an Eulerian tour \mathcal{E} as follows. Choose an edge e_1 with $\text{init}(e_1) = v$. Then continue to choose any edge possible to continue the tour, except we never choose an edge f of T unless we have to, i.e., unless it's the only remaining edge exiting the vertex at which we stand. Then we never get stuck until all edges are used, so we have constructed an Eulerian tour \mathcal{E} . Moreover, the set of last exits of \mathcal{E} from vertices $u \neq v$ of D coincides with the set of edges of the oriented tree T .

Proof of Claim #2. Since D is balanced, the only way to get stuck is to end up at v with no further exits available, but with an edge still unused. Suppose this is the case. At least one unused edge must be a last exit edge, i.e., an edge of T [why?]. Let u be a vertex of T closest to v in T such that the unique edge f of T with $\text{init}(f) = u$ is not in the tour. Let $y = \text{fin}(f)$. Suppose $y \neq v$. Since we enter y as often as we leave it, we don't use the last exit from y . Thus $y = v$. But then we can leave v , a contradiction. This proves Claim #2.

We have shown that every Eulerian tour \mathcal{E} beginning with the edge e has associated with it a “last exit” oriented subtree $T = T(\mathcal{E})$ with root $v = \text{init}(e)$. Conversely, given an oriented subtree T with root v , we can obtain all Eulerian tours \mathcal{E} beginning with e and satisfying $T = T(\mathcal{E})$ by choosing for each vertex $u \neq v$ the order in which the edges from u , except the edge of T , appear in \mathcal{E} ; as well as choosing the order in which all the edges from v except for e appear in \mathcal{E} . Thus for each vertex u we have $(\text{outdeg}(u) - 1)!$ choices, so for each T we have $\prod_u (\text{outdeg}(u) - 1)!$ choices. Since there are $\tau(D, v)$ choices for T , the proof is complete. \square

10.3 Corollary. *Let D be a connected balanced digraph, and let v be a vertex of D . Then the number $\tau(D, v)$ of oriented subtrees with root v is independent of v .*

Proof. Let e be an edge with initial vertex v . By equation (10.1), we need to show that the number $\epsilon(D, e)$ of Eulerian tours beginning with e is independent of e . But $e_1 e_2 \cdots e_q$ is an Eulerian tour if and only if $e_i e_{i+1} \cdots e_q e_1 e_2 \cdots e_{i-1}$ is also an Eulerian tour, and the proof follows [why?]. \square

What we obviously need to do next is find a formula for $\tau(D, v)$. This result turns out to be very similar to the Matrix-Tree Theorem, and indeed we will show (Example 10.6) that the Matrix-Tree Theorem is a simple corollary to Theorem 10.4.

10.4 Theorem. *Let D be a connected digraph with vertex set $V = \{v_1, \dots, v_p\}$ and with l_i loops at vertex v_i . Let $\mathbf{L}(D)$ be the $p \times p$ matrix defined by*

$$\mathbf{L}_{ij} = \begin{cases} -m_{ij}, & \text{if } i \neq j \text{ and there are } m_{ij} \text{ edges with} \\ & \text{initial vertex } v_i \text{ and final vertex } v_j \\ \text{outdeg}(v_i) - l_i, & \text{if } i = j. \end{cases}$$

(Thus \mathbf{L} is the directed analogue of the laplacian matrix of an undirected graph.) Let \mathbf{L}_0 denote \mathbf{L} with the last row and column deleted. Then

$$\det \mathbf{L}_0 = \tau(D, v_p). \quad (10.2)$$

NOTE. If we remove the i th row and column from \mathbf{L} instead of the last row and column, then equation (10.2) still holds with v_p replaced with v_i .

Proof (sketch). Induction on q , the number of edges of D . The fewest number of edges which D can have is $p - 1$ (since D is connected). Suppose then that D has $p - 1$ edges, so that as an undirected graph D is a tree. If D is not an oriented tree with root v_p , then some vertex $v_i \neq v_p$ of D has outdegree 0 [why?]. Then \mathbf{L}_0 has a zero row, so $\det \mathbf{L}_0 = 0 = \tau(D, v_p)$. If on the other hand D is an oriented tree with root v_p , then an argument like that used to prove Lemma 9.7 (in the case when S is the set of edges of a spanning tree) shows that $\det \mathbf{L}_0 = 1 = \tau(D, v_p)$.

Now assume that D has $q > p - 1$ edges, and assume the theorem for digraphs with at most $q - 1$ edges. We may assume that no edge f of D

has initial vertex v_p , since such an edge belongs to no oriented tree with root v_p and also makes no contribution to \mathbf{L}_0 . It then follows, since D has at least p edges, that there exists a vertex $u \neq v_p$ of D of outdegree at least two. Let e be an edge with $\text{init}(e) = u$. Let D_1 be D with the edge e removed. Let D_2 be D with all edges e' removed such that $\text{init}(e) = \text{init}(e')$ and $e' \neq e$. (Note that D_2 is strictly smaller than D since $\text{outdeg}(u) \geq 2$.) By induction, we have $\det \mathbf{L}_0(D_1) = \tau(D_1, v_p)$ and $\det \mathbf{L}_0(D_2) = \tau(D_2, v_p)$. Clearly $\tau(D, v_p) = \tau(D_1, v_p) + \tau(D_2, v_p)$, since in an oriented tree T with root v_p , there is exactly one edge whose initial vertex coincides with that of e . On the other hand, it follows immediately from the multilinearity of the determinant [why?] that

$$\det \mathbf{L}_0(D) = \det \mathbf{L}_0(D_1) + \det \mathbf{L}_0(D_2).$$

From this the proof follows by induction. \square

10.5 Corollary. *Let D be a connected balanced digraph with vertex set $V = \{v_1, \dots, v_p\}$. Let e be an edge of D . Then the number $\epsilon(D, e)$ of Eulerian tours of D with first edge e is given by*

$$\epsilon(D, e) = (\det \mathbf{L}_0(D)) \prod_{u \in V} (\text{outdeg}(u) - 1)!.$$


Equivalently (since D is balanced, so Lemma 9.9 applies), if $\mathbf{L}(D)$ has eigenvalues μ_1, \dots, μ_p with $\mu_p = 0$, then

$$\epsilon(D, e) = \frac{1}{p} \mu_1 \cdots \mu_{p-1} \prod_{u \in V} (\text{outdeg}(u) - 1)!.$$

Proof. Combine Theorems 10.2 and 10.4. \square

10.6 Example (the Matrix-Tree Theorem revisited). Let $G = (V, E, \varphi)$ be a connected loopless undirected graph. Let $\widehat{G} = (V, \widehat{E}, \widehat{\varphi})$ be the digraph obtained from G by replacing each edge e of G , where $\varphi(e) = \{u, v\}$, with a pair e' and e'' of directed edges satisfying $\widehat{\varphi}(e') = (u, v)$ and $\widehat{\varphi}(e'') = (v, u)$. Clearly \widehat{G} is balanced and connected. Choose a vertex v of G . There is an obvious one-to-one correspondence between spanning trees T of G and oriented spanning trees \widehat{T} of \widehat{G} with root v , namely, direct each edge of T toward v . Moreover, $\mathbf{L}(G) = \mathbf{L}(\widehat{G})$ [why?]. Hence the Matrix-Tree Theorem is an immediate consequence of Theorem 10.4.

10.7 Example (the efficient mail carrier). A mail carrier has an itinerary of city blocks to which he (or she) must deliver mail. He wants to accomplish this by walking along each block twice, once in each direction, thus passing along houses on each side of the street. He also wants to end up where he started, which is where his car is parked. The blocks form the edges of a graph G , whose vertices are the intersections. The mail carrier wants simply to walk along an Eulerian tour in the digraph \widehat{G} of the previous example. Making the plausible assumption that the graph is connected, not only does an Eulerian tour always exist, but we can tell the mail carrier how many there are. Thus he will know how many different routes he can take to avoid boredom. For instance, suppose G is the 3×3 grid illustrated below.



This graph has 192 spanning trees. Hence the number of mail carrier routes beginning with a fixed edge (in a given direction) is $192 \cdot 1!^4 2!^4 3! = 18432$. The total number of routes is thus 18432 times twice the number of edges [why?], viz., $18432 \times 24 = 442368$. Assuming the mail carrier delivered mail 250 days a year, it would be 1769 years before he would have to repeat a route!


10.8 Example (binary de Bruijn sequences). A *binary sequence* is just a sequence of 0's and 1's. A *binary de Bruijn sequence* of degree n is a binary sequence $A = a_1 a_2 \cdots a_{2^n}$ such that every binary sequence $b_1 \cdots b_n$ of length n occurs exactly once as a “circular factor” of A , i.e., as a sequence $a_i a_{i+1} \cdots a_{i+n-1}$, where the subscripts are taken modulo 2^n if necessary. For instance, some circular factors of the sequence $abcdefg$ are a , $bcde$, $fgab$, and $defga$. Note that there are exactly 2^n binary sequences of length n , so the only possible length of a binary de Bruijn sequence of degree n is 2^n [why?]. Clearly any cyclic shift $a_i a_{i+1} \cdots a_{2^n} a_1 a_2 \cdots a_{i-1}$ of a binary de Bruijn sequence $a_1 a_2 \cdots a_{2^n}$ is also a binary de Bruijn sequence, and we call two such sequences *equivalent*. This relation of equivalence is obviously an equivalence relation, and every equivalence class contains exactly one sequence beginning with n 0's [why?]. Up to equivalence, there is one binary de Bruijn sequence of degree two, namely, 0011. It's easy to check that there are two inequivalent binary de Bruijn sequences of degree three, namely, 00010111 and 00011101.

However, it's not clear at this point whether binary de Bruijn sequences exist for all n . By a clever application of Theorems 10.2 and 10.4, we will not only show that such sequences exist for all positive integers n , but we will also count the number of them. It turns out that there are *lots* of them. For instance, the number of inequivalent binary de Bruijn sequences of degree eight is equal to

$$1329227995784915872903807060280344576,$$

as the reader can easily check by writing down all these sequences. De Bruijn sequences have a number of interesting applications to the design of switching networks and related topics.

Our method of enumerating binary de Bruijn sequences will be to set up a correspondence between them and Eulerian tours in a certain directed graph D_n , the *de Bruijn graph* of degree n . The graph D_n has 2^{n-1} vertices, which we will take to consist of the 2^{n-1} binary sequences of length $n-1$. A pair $(a_1a_2 \cdots a_{n-1}, b_1b_2 \cdots b_{n-1})$ of vertices forms an edge of D_n if and only if $a_2a_3 \cdots a_{n-1} = b_1b_2 \cdots b_{n-2}$, i.e., e is an edge if the last $n-2$ terms of $\text{init}(e)$ agree with the first $n-2$ terms of $\text{fin}(e)$. Thus every vertex has indegree two and outdegree two [why?], so D_n is balanced. The number of edges of D_n is 2^n . Moreover, it's easy to see that D_n is connected (see Lemma 10.9). The graphs D_3 and D_4 look as follows:



Suppose that $\mathcal{E} = e_1e_2 \cdots e_{2^n}$ is an Eulerian tour in D_n . If $\text{fin}(e_i)$ is the binary sequence $a_{i,1}a_{i,2} \cdots a_{i,n-1}$, then replace e_i in \mathcal{E} by the last bit $a_{i,n-1}$.

For instance, the Eulerian tour (where we simply write the vertices)

000, 000, 001, 010, 101, 011, 111, 111, 110, 101, 010, 100, 001, 011, 110, 100, 000

corresponds to the de Bruijn sequence 0101111010011000 (the last bits of the vertices above, excluding the first vertex 000). It is easy to see that the resulting sequence $\beta(\mathcal{E}) = a_{1,n-1}a_{2,n-1} \cdots a_{2^n,n-1}$ is a binary de Bruijn sequence, and conversely every binary de Bruijn sequence arises in this way. In particular, since D_n is balanced and connected there exists at least one binary de Bruijn sequence. In order to count the total number of such sequences, we need to compute $\det \mathbf{L}(D_n)$. One way to do this is by a clever but messy sequence of elementary row and column operations which transforms the determinant into triangular form. We will give instead an elegant computation of the eigenvalues of $\mathbf{L}(D_n)$ based on the following simple lemma.

10.9 Lemma. *Let u and v be any two vertices of D_n . Then there is a unique (directed) walk from u to v of length $n - 1$.*

Proof. Suppose $u = a_1a_2 \cdots a_{n-1}$ and $v = b_1b_2 \cdots b_{n-1}$. Then the unique path of length $n - 1$ from u to v has vertices

$$a_1a_2 \cdots a_{n-1}, a_2a_3 \cdots a_{n-1}b_1, a_3a_4 \cdots a_{n-1}b_1b_2, \dots, \\ a_{n-1}b_1 \cdots b_{n-2}, b_1b_2 \cdots b_{n-1}.$$

□

10.10 Theorem. *The eigenvalues of $\mathbf{L}(D_n)$ are 0 (with multiplicity one) and 2 (with multiplicity $2^{n-1} - 1$).*

Proof. Let $\mathbf{A}(D_n)$ denote the directed adjacency matrix of D_n , i.e., the rows and columns are indexed by the vertices, with

$$\mathbf{A}_{uv} = \begin{cases} 1, & \text{if } (u, v) \text{ is an edge} \\ 0, & \text{otherwise.} \end{cases}$$

Now Lemma 10.9 is equivalent to the assertion that $\mathbf{A}^{n-1} = J$, the $2^{n-1} \times 2^{n-1}$ matrix of all 1's [why?]. If the eigenvalues of \mathbf{A} are $\lambda_1, \dots, \lambda_{2^{n-1}}$, then the eigenvalues of $J = \mathbf{A}^{n-1}$ are $\lambda_1^{n-1}, \dots, \lambda_{2^{n-1}}^{n-1}$. By Lemma 1.4, the eigenvalues of J are 2^{n-1} (once) and 0 ($2^{n-1} - 1$ times). Hence the eigenvalues of \mathbf{A} are 2ζ (once, where ζ is an $(n - 1)$ -st root of unity to be determined), and 0

($2^{n-1} - 1$ times). Since the trace of \mathbf{A} is 2, it follows that $\zeta = 1$, and we have found all the eigenvalues of \mathbf{A} .

Now $\mathbf{L}(D_n) = 2I - \mathbf{A}(D_n)$ [why?]. Hence the eigenvalues of \mathbf{L} are $2 - \lambda_1, \dots, 2 - \lambda_{2^{n-1}}$, and the proof follows from the above determination of $\lambda_1, \dots, \lambda_{2^{n-1}}$. \square

10.11 Corollary. *The number $B_0(n)$ of binary de Bruijn sequences of degree n beginning with n 0's is equal to $2^{2^{n-1}-n}$. The total number $B(n)$ of binary de Bruijn sequences of degree n is equal to $2^{2^{n-1}}$.*

Proof. By the above discussion, $B_0(n)$ is the number of Eulerian tours in D_n whose first edge is the loop at vertex $00 \cdots 0$. Moreover, the outdegree of every vertex of D_n is two. Hence by Corollary 10.5 and Theorem 10.10 we have

$$B_0(n) = \frac{1}{2^{n-1}} 2^{2^{n-1}-1} = 2^{2^{n-1}-n}.$$

Finally, $B(n)$ is obtained from $B_0(n)$ by multiplying by the number 2^n of edges, and the proof follows. \square

Note that the total number of binary sequences of length 2^n is $N = 2^{2^n}$. By the previous corollary, the number of these which are de Bruijn sequences is just \sqrt{N} . This suggests the following problem, which remained open until 2009. Let \mathcal{A}_n be the set of all binary sequences of length 2^n . Let \mathcal{B}_n be the set of binary de Bruijn sequences of degree n . Find an explicit bijection

$$\psi: \mathcal{B}_n \times \mathcal{B}_n \rightarrow \mathcal{A}_n, \tag{10.3}$$

thereby giving a combinatorial proof of Corollary 10.11.

Notes for Chapter 10

The characterization of Eulerian digraphs given by Theorem 10.1 is a result of I. J. Good [46], while the fundamental connection between oriented subtrees and Eulerian tours in a balanced digraph that was used to prove Theorem 10.2 was shown by T. van Aardenne-Ehrenfest and N. G. de Bruijn [1, Thm. 5a]. This result is sometimes called the BEST Theorem, after de Bruijn, van Aardenne-Ehrenfest, Smith, and Tutte. However, Smith and Tutte were not involved in the original discovery. (In [101] Smith and Tutte give a determinantal formula for the number of Eulerian tours in a

special class of balanced digraphs. Van Aardenne-Ehrenfest and de Bruijn refer to the paper of Smith and Tutte in a footnote added in proof.) The determinantal formula for the number of oriented subtrees of a directed graph (Theorem 10.4) is due to Tutte [118, Thm. 3.6].

De Bruijn sequences are named from the paper [15] of de Bruijn, where they are enumerated in the binary case. However, it was discovered by R. Stanley in 1975 that this work had been done earlier by C. Flye Sainte-Marie [36] in 1894, as reported by de Bruijn [17]. The generalization to d -ary de Bruijn sequences (Exercise 10.2) is to T. van Ardenne-Ehrenfest and de Bruijn [1]. Some recent work in this area appears in a special issue [111] of *Discrete Mathematics*. Some entertaining applications to magic are given by P. Diaconis and R. Graham [30, Chs. 2–4]. The bijection ψ of equation (10.3) is due to H. Bidkhori and S. Kishore. [9].

Chapter 11


Cycles, bonds, and electrical networks

11.1 The cycle space and bond space

In this chapter we will deal with some interesting linear algebra related to the structure of a directed graph. Let $D = (V, E)$ be a digraph. A function $f: E \rightarrow \mathbb{R}$ is called a *circulation* if for every vertex $v \in V$, we have


$$\sum_{\substack{e \in E \\ \text{init}(e)=v}} f(e) = \sum_{\substack{e \in E \\ \text{fin}(e)=v}} f(e). \quad (11.1)$$

Thus if we think of the edges as pipes and f as measuring the flow (quantity per unit of time) of some commodity (such as oil) through the pipes in the specified direction (so that a negative value of $f(e)$ means a flow of $|f(e)|$ in the direction opposite the direction of e), then equation (11.1) simply says that the amount flowing into each vertex equals the amount flowing out. In other words, the flow is *conservative*. The figure below illustrates a circulation in a digraph D .



Let $\mathcal{C} = \mathcal{C}_D$ denote the set of all circulations on D . Clearly if $f, g \in \mathcal{C}$ and $\alpha, \beta \in \mathbb{R}$ then $\alpha f + \beta g \in \mathcal{C}$. Hence \mathcal{C} is a (real) vector space, called the *cycle space* of D . Thus if $q = \#E$, then \mathcal{C}_D is a subspace of the q -dimensional vector space \mathbb{R}^E of all functions $f: E \rightarrow \mathbb{R}$.


What do circulations have to do with something “circulating,” and what does the cycle space have to do with actual cycles? To see this, define a *circuit* or *elementary cycle* in D to be a set of edges of a closed walk, *ignoring the direction of the arrows*, with no repeated vertices except the first and last. Suppose that C has been assigned an orientation (direction of travel) \mathfrak{o} . (Note that this meaning of orientation is not the same as that appearing in Definition 9.5.)



Define a function $f_C: E \rightarrow \mathbb{R}$ (which also depends on the orientation \mathfrak{o} , though we suppress it from the notation) by

$$f_C(e) = \begin{cases} 1, & \text{if } e \in C \text{ and } e \text{ agrees with } \mathfrak{o} \\ -1, & \text{if } e \in C \text{ and } e \text{ is opposite to } \mathfrak{o} \\ 0, & \text{otherwise.} \end{cases}$$

It is easy to see that f_C is a circulation. Later we will see that the circulations f_C span the cycle space \mathcal{C} , explaining the terminology “circulation” and “cycle space.” The figure below shows a circuit C with an orientation \mathfrak{o} , and the corresponding circulation f_C .



Given a function $\varphi: V \rightarrow \mathbb{R}$, called a *potential* on D , define a new function $\delta\varphi: E \rightarrow \mathbb{R}$, called the *coboundary*¹ of φ , by

$$\delta\varphi(e) = \varphi(v) - \varphi(u), \text{ if } u = \text{init}(e) \text{ and } v = \text{fin}(e).$$

Figure 11.1 shows a digraph D with the value $\varphi(v)$ of some function $\varphi: V \rightarrow \mathbb{R}$ indicated at each vertex v , and the corresponding values $\delta\varphi(e)$ shown at each edge e .

One should regard δ as an operator which takes an element φ of the vector space \mathbb{R}^V of all functions $V \rightarrow \mathbb{R}$ and produces an element of the vector space \mathbb{R}^E of all functions $E \rightarrow \mathbb{R}$. It is immediate from the definition of δ that δ is *linear*, i.e.,

$$\delta(a\varphi_1 + b\varphi_2) = a \cdot \delta\varphi_1 + b \cdot \delta\varphi_2,$$

for all $\varphi_1, \varphi_2 \in \mathbb{R}^V$ and $a, b \in \mathbb{R}$. Thus δ is simply a linear transformation $\delta: \mathbb{R}^V \rightarrow \mathbb{R}^E$ between two finite-dimensional vector spaces.

A function $g: E \rightarrow \mathbb{R}$ is called a *potential difference* on D if $g = \delta\varphi$ for some $\varphi: V \rightarrow \mathbb{R}$. (Later we will see the connection with electrical networks that accounts for the terminology “potential difference.”) Let $\mathcal{B} = \mathcal{B}_D$ be the set of all potential differences on D . Thus \mathcal{B} is just the image of the linear transformation δ and is hence a real vector space, called the *bond space* of D .

Let us explain the reason behind the terminology “bond space.” A *bond* in a digraph D is a set B of edges such that (a) removing B from D disconnects





Figure 11.1: A function (potential) and its coboundary

¹The term “coboundary” arises from algebraic topology, but we will not explain the connection here.

some (undirected) component of D (that is, removing B creates a digraph which has more connected components, as an undirected graph, than D), and (b) no proper subset of B has this property. A subset of edges satisfying (a) is called a *cutset*, so a bond is just a minimal cutset. Suppose, for example, that D is given as follows (with no arrows drawn since they are irrelevant to the definition of bond):



Then the bonds are the six subsets $ab, de, acd, bce, ace, bcd$.

Let B be a bond. Suppose B disconnects the component (V', E') into two pieces (a bond always disconnects some component into exactly two pieces [why?]) with vertex set S in one piece and \bar{S} in the other. Thus $S \cup \bar{S} = V'$ and $S \cap \bar{S} = \emptyset$. Define

$$[S, \bar{S}] = \{e \in E : \text{exactly one vertex of } e \text{ lies in } S \text{ and one lies in } \bar{S}\}.$$

Clearly $B = [S, \bar{S}]$. It is often convenient to use the notation $[S, \bar{S}]$ for a bond.

Given a bond $B = [S, \bar{S}]$ of D , define a function $g_B: E \rightarrow \mathbb{R}$ by

$$g_B(e) = \begin{cases} 1, & \text{if } \text{init}(e) \in \bar{S}, \text{fin}(e) \in S \\ -1, & \text{if } \text{init}(e) \in S, \text{fin}(e) \in \bar{S} \\ 0, & \text{otherwise.} \end{cases}$$


Note that g_B really depends not just on B , but on whether we write B as $[S, \bar{S}]$ or $[\bar{S}, S]$. Writing B in the reverse way simply changes the sign of g_B . Whenever we deal with g_B we will assume that some choice $B = [S, \bar{S}]$ has been made.

Now note that $g_B = \delta\varphi$, where

$$\varphi(v) = \begin{cases} 1, & \text{if } v \in S \\ 0, & \text{if } v \notin S. \end{cases}$$

Hence $g_B \in \mathcal{B}$, the bond space of D . We will later see that \mathcal{B} is in fact spanned by the functions g_B , explaining the terminology “bond space.”


11.1 Example. In the digraph below, open (white) vertices indicate an element of S and closed (black) vertices an element of \bar{S} for a certain bond $B = [S, \bar{S}]$. The elements of B are drawn with solid lines. The edges are labelled by the values of g_B , and the vertices by the function φ for which $g_B = \delta\varphi$.



Recall that in Definition 9.5 we defined the incidence matrix $\mathbf{M}(G)$ of a loopless undirected graph G with respect to an orientation \mathfrak{o} . We may just as well think of G together with its orientation \mathfrak{o} as a directed graph. We also will allow loops. Thus if $D = (V, E)$ is any (finite) digraph, define the *incidence matrix* $\mathbf{M} = \mathbf{M}(D)$ to be the $p \times q$ matrix whose rows are indexed by V and columns by E , as follows. The entry in row $v \in V$ and column $e \in E$ is denoted $m_v(e)$ and is given by

$$m_v(e) = \begin{cases} -1, & \text{if } v = \text{init}(e) \text{ and } e \text{ is not a loop} \\ 1, & \text{if } v = \text{fin}(e) \text{ and } e \text{ is not a loop} \\ 0, & \text{otherwise.} \end{cases}$$

For instance, if D is given by



then

$$\mathbf{M}(D) = \begin{bmatrix} 1 & 1 & -1 & 0 & 0 \\ -1 & -1 & 0 & 1 & 0 \\ 0 & 0 & 1 & -1 & 0 \end{bmatrix}.$$

11.2 Theorem. *The row space of $\mathbf{M}(D)$ is the bond space \mathcal{B}_D . Equivalently, the functions $m_v: E \rightarrow \mathbb{R}$, where v ranges over all vertices of D , span \mathcal{B}_D .*

Proof. Let $g = \delta\varphi$ be a potential difference on D , so

$$\begin{aligned} g(e) &= \varphi(\text{fin}(e)) - \varphi(\text{init}(e)) \\ &= \sum_{v \in V} \varphi(v) m_v(e). \end{aligned}$$

Thus $g = \sum_{v \in V} \varphi(v) m_v$, so g belongs to the row space of \mathbf{M} .

Conversely, if $g = \sum_{v \in V} \psi(v) m_v$ is in the row space of \mathbf{M} , where $\psi: V \rightarrow \mathbb{R}$, then $g = \delta\psi \in \mathcal{B}$. \square

We now define a scalar product (or inner product) on the space \mathbb{R}^E by

$$\langle f, g \rangle = \sum_{e \in E} f(e)g(e),$$

for any $f, g \in \mathbb{R}^E$. If we think of the numbers $f(e)$ and $g(e)$ as the coordinates of f and g with respect to the basis E , then $\langle f, g \rangle$ is just the usual dot product of f and g . Because we have a scalar product, we have a notion of what it means for f and g to be *orthogonal*, viz., $\langle f, g \rangle = 0$. If \mathcal{V} is any subspace of \mathbb{R}^E , then define the *orthogonal complement* \mathcal{V}^\perp of \mathcal{V} by

$$\mathcal{V}^\perp = \{f \in \mathbb{R}^E : \langle f, g \rangle = 0 \text{ for all } g \in \mathcal{V}\}.$$

Recall from linear algebra that

$$\dim \mathcal{V} + \dim \mathcal{V}^\perp = \dim \mathbb{R}^E = \#E. \quad (11.2)$$

Furthermore, $(\mathcal{V}^\perp)^\perp = \mathcal{V}$. Let us also note that since we are working over \mathbb{R} , we have $\mathcal{V} \cap \mathcal{V}^\perp = \{0\}$. Thus $\mathbb{R}^E = \mathcal{V} \oplus \mathcal{V}^\perp$ (direct sum).

Intuitively there is a kind of “duality” between elementary cycles and bonds. Cycles “hold vertices together,” while bonds “tear them apart.” The precise statement of this duality is given by the next result.

11.3 Theorem. *The cycle and bond spaces of D are related by $\mathcal{C} = \mathcal{B}^\perp$. (Equivalently, $\mathcal{B} = \mathcal{C}^\perp$.)*

Proof. Let $f: E \rightarrow \mathbb{R}$. Then f is a circulation if and only if

$$\sum_{e \in E} m_v(e) f(e) = 0$$

for all $v \in V$ [why?]. But this is exactly the condition that $f \in \mathcal{B}^\perp$. \square

11.2 Bases for the cycle space and bond space

We want to examine the incidence matrix $\mathbf{M}(D)$ in more detail. In particular, we would like to determine which rows and columns of $\mathbf{M}(D)$ are linearly independent, and which span the row and column spaces. As a corollary, we will determine the dimension of the spaces \mathcal{B} and \mathcal{C} . We begin by defining the *support* $\|f\|$ of $f: E \rightarrow \mathbb{R}$ to be the set of edges $e \in E$ for which $f(e) \neq 0$.

11.4 Lemma. *If $0 \neq f \in \mathcal{C}$, then $\|f\|$ contains an undirected circuit.*

Proof. If not, then $\|f\|$ has a vertex of degree one [why?], which is clearly impossible. \square

11.5 Lemma. *If $0 \neq g \in \mathcal{B}$, then $\|g\|$ contains a bond.*

Proof. Let $0 \neq g \in \mathcal{B}$, so $g = \delta\varphi$ for some $\varphi: V \rightarrow \mathbb{R}$. Choose a vertex v which is incident to an edge of $\|g\|$, and set

$$U = \{u \in V: \varphi(u) = \varphi(v)\}.$$

Let $\bar{U} = V - U$. Note that $\bar{U} \neq \emptyset$, since otherwise φ is constant so $g = 0$. Since $g(e) \neq 0$ for all $e \in [U, \bar{U}]$ [why?], we have that $\|g\|$ contains the cutset $[U, \bar{U}]$. Since a bond is by definition a minimal cutset, it follows that $\|g\|$ contains a bond. \square

A matrix \mathbf{B} is called a *basis matrix* of \mathcal{B} if the rows of \mathbf{B} form a basis for \mathcal{B} . Similarly define a basis matrix \mathbf{C} of \mathcal{C} .

Recall the notation of Theorem 9.4: let A be a matrix with at least as many columns as rows, whose columns are indexed by the elements of a set T . If $S \subseteq T$, then $A[S]$ denotes the submatrix of A consisting of the columns indexed by the elements of S . In particular, $A[e]$ (short for $A[\{e\}]$) denotes the column of A indexed by e . We come to our first significant result about bases for the vector spaces \mathcal{B} and \mathcal{C} .

11.6 Theorem. Let \mathbf{B} be a basis matrix of \mathcal{B} , and \mathbf{C} a basis matrix of \mathcal{C} . (Thus the columns of \mathbf{B} and \mathbf{C} are indexed by the edges $e \in E$ of D .) Let $S \subseteq E$, Then:

- (i) The columns of $\mathbf{B}[S]$ are linearly independent if and only if S is acyclic (i.e., contains no circuit as an undirected graph).
- (ii) The columns of $\mathbf{C}[S]$ are linearly independent if and only if S contains no bond.

Proof. The columns of $\mathbf{B}[S]$ are linearly dependent if and only if there exists a function $f: E \rightarrow \mathbb{R}$ such that

$$f(e) \neq 0 \text{ for some } e \in S$$

$$f(e) = 0 \text{ for all } e \notin S$$

$$\sum_{e \in E} f(e) \mathbf{B}[e] = \mathbf{0}, \text{ the column vector of 0's.} \quad (11.3)$$

The last condition is equivalent to $\langle f, m_v \rangle = 0$ for all $v \in V$, i.e., f is a circulation. Thus the columns of $\mathbf{B}[S]$ are linearly dependent if and only if there exists a nonzero circulation f such that $\|f\| \subseteq S$. By Lemma 11.4, $\|f\|$ (and therefore S) contains a circuit. Conversely, if S contains a circuit C then $0 \neq f_C \in \mathcal{C}$ and $\|f_C\| = C \subseteq S$, so f_C defines a linear dependence relation (11.3) among the columns. Hence the columns of $\mathbf{B}[S]$ are linearly independent if and only if S is acyclic, proving (i). (Part (i) can also be deduced from Lemma 9.7.)

The proof of (ii) is similar and is left as an exercise. \square

11.7 Corollary. Let $D = (V, E)$ be a digraph with p vertices, q edges, and k connected components (as an undirected graph). Then

$$\begin{aligned} \dim \mathcal{B} &= p - k \\ \dim \mathcal{C} &= q - p + k. \end{aligned}$$

Proof. For any matrix X , the rank of X is equal to the maximum number of linearly independent columns. Now let \mathbf{B} be a basis matrix of \mathcal{B} . By Theorem 11.6(i), the rank of \mathbf{B} is then the maximum size (number of elements) of an acyclic subset of E . In each connected component D_i of D , the largest

acyclic subsets are the spanning trees, whose number of edges is $p(D_i) - 1$, where $p(D_i)$ is the number of vertices of D_i . Hence

$$\begin{aligned} \text{rank } \mathcal{B} &= \sum_{i=1}^k (p(D_i) - 1) \\ &= p - k. \end{aligned}$$

Since $\dim \mathcal{B} + \dim \mathcal{C} = \dim \mathbb{R}^E = q$ by equation (11.2) and Theorem 11.3, we have

$$\dim \mathcal{C} = q - (p - k) = q - p + k.$$

(It is also possible to determine $\dim \mathcal{C}$ by a direct argument similar to our determination of $\dim \mathcal{B}$.) \square

The number $q - p + k$ (which should be thought of as the number of independent cycles in D) is called the *cyclomatic number* of D (or of its undirected version G , since the direction of the edges have no effect).

Our next goal is to describe explicit bases of \mathcal{C} and \mathcal{B} . Recall that a *forest* is an undirected graph without circuits, or equivalently, a disjoint union of trees. We extend the definition of forest to directed graphs by ignoring the arrows, i.e., a directed graph is a forest if it has no circuits as an undirected graph. Equivalently [why?], $\dim \mathcal{C} = 0$.

Pick a maximal forest T of $D = (V, E)$. Thus T restricted to each component of D is a spanning tree. If e is an edge of D not in T , then it is easy to see that $T \cup e$ contains a unique circuit C_e .


11.8 Theorem. *Let T be as above. Then the set S of circulations f_{C_e} , as e ranges over all edges of D not in T , is a basis for the cycle space \mathcal{C} .*

Proof. The circulations f_{C_e} are linearly independent, since for each $e \in E(D) - E(T)$ only f_{C_e} doesn't vanish on e . Moreover,

$$\#S = \#E(D) - \#E(T) = q - p + k = \dim \mathcal{C},$$

so S is a basis. \square

11.9 Example. Let D be the digraph shown below, with the edges a, b, c of T shown by dotted lines.



Orient each circuit C_t in the direction of the added edge, i.e., $f_{C_t}(t) = 1$. Then the basis matrix \mathbf{C} of \mathcal{C} corresponding to the basis $f_{C_d}, f_{C_e}, f_{C_f}$ is given by

$$\mathbf{C} = \begin{bmatrix} 0 & -1 & -1 & 1 & 0 & 0 \\ -1 & -1 & -1 & 0 & 1 & 0 \\ 0 & 0 & -1 & 0 & 0 & 1 \end{bmatrix}. \quad (11.4)$$

We next want to find a basis for the bond space \mathcal{B} analogous to that of Theorem 11.8.

11.10 Lemma. *Let T be a maximal forest of $D = (V, E)$. Let $T^* = D - E(T)$ (the digraph obtained from D by removing the edges of T), called a cotree if D is connected. Let e be an edge of T . Then $E(T^*) \cup e$ contains a unique bond.*

Proof. Removing $E(T^*)$ from D leaves a maximal forest T , so removing one further edge e disconnects some component of D . Hence $E(T^*) \cup e$ contains a bond B . It remains to show that B is unique. Removing e from T breaks some component of T into two connected graphs T_1 and T_2 with vertex sets S and \bar{S} . It follows [why?] that we must have $B = [S, \bar{S}]$, so B is unique. \square

Let T be a maximal forest of the digraph D , and let e be an edge of T . By the previous lemma, $E(T^*) \cup e$ contains a unique bond B_e . Let g_{B_e} be the corresponding element of the bond space \mathcal{B} , chosen for definiteness so that $g_{B_e}(e) = 1$.

11.11 Theorem. *The set of functions g_{B_e} , as e ranges over all edges of T , is a basis for the bond space \mathcal{B} .*

Proof. The functions g_{B_e} are linearly independent, since only g_{B_e} is nonzero on $e \in E(T)$. Since

$$\#E(T) = p - k = \dim \mathcal{B},$$

it follows that the g_{B_e} 's are a basis for \mathcal{B} . \square

11.12 Example. Let D and T be as in the previous diagram. Thus a basis for \mathcal{B} is given by the functions $g_{B_a}, g_{B_b}, g_{B_c}$. The corresponding basis matrix is given by

$$\mathbf{B} = \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 & 1 & 1 \end{bmatrix}.$$

Note that the rows of \mathbf{B} are orthogonal to the rows of the matrix \mathbf{C} of equation (11.4), in accordance with Theorem 11.3. Equivalently, $\mathbf{B}\mathbf{C}^t = \mathbf{0}$, the 3×3 zero matrix. (In general, $\mathbf{B}\mathbf{C}^t$ will have $q - p + k$ rows and $p - k$ columns. Here it is just a coincidence that these two numbers are equal.)

The basis matrices \mathbf{C}_T and \mathbf{B}_T of \mathcal{C} and \mathcal{B} obtained from a maximal forest T have an important property. A real matrix $m \times n$ matrix A with $m \leq n$ is said to be *unimodular* if every $m \times m$ submatrix has determinant 0, 1, or -1 . For instance, the adjacency matrix $\mathbf{M}(D)$ of a digraph D is unimodular, as proved in Lemma 9.7 (by showing that the expansion of the determinant of a full submatrix has at most one nonzero term).

11.13 Theorem. *Let T be a maximal forest of D . Then the basis matrices \mathbf{C}_T of \mathcal{C} and \mathbf{B}_T of \mathcal{B} are unimodular.*

Proof. First consider the case \mathbf{C}_T . Let \mathbf{P} be a full submatrix of \mathbf{C} (so \mathbf{P} has $q - p + k$ rows and columns). Assume $\det \mathbf{P} \neq 0$. We need to show $\det \mathbf{P} = \pm 1$. Since $\det \mathbf{P} \neq 0$, it follows from Theorem 11.6(ii) that $\mathbf{P} = \mathbf{C}_T[T_1^*]$ for the complement T_1^* of some maximal forest T_1 . Note that the rows of the matrix $\mathbf{C}_T[T_1^*]$ are indexed by T^* and the columns by T_1^* . Similarly the rows of the basis matrix \mathbf{C}_{T_1} are indexed by T_1^* and the columns by E (the set of all edges of D). Hence it makes sense to define the matrix product

$$\mathbf{Z} = \mathbf{C}_T[T_1^*]\mathbf{C}_{T_1},$$

a matrix whose rows are indexed by T^* and columns by E .

Note that the matrix \mathbf{Z} is a basis matrix for the cycle space \mathcal{C} since its rows are linear combinations of the rows of the basis matrix $\mathbf{C}_{T_1}^*$, and it has full rank since the matrix $\mathbf{C}_T[T_1^*]$ is invertible. Now $\mathbf{C}_{T_1}[T_1^*] = I_{T_1^*}$ (the identity matrix indexed by T_1^*), so $\mathbf{Z}[T_1^*] = \mathbf{C}_T[T_1^*]$. Thus \mathbf{Z} agrees with the basis matrix \mathbf{C}_T in columns T_1^* . Hence the rows of $\mathbf{Z} - \mathbf{C}_T$ are circulations supported on a subset of T_1 . Since T_1 is acyclic, it follows from Lemma 11.4 that the only such circulation is identically 0, so $\mathbf{Z} = \mathbf{C}_T$.

We have just shown that

$$\mathbf{C}_T[T_1^*]\mathbf{C}_{T_1} = \mathbf{C}_T.$$

Restricting both sides to T^* , we obtain

$$\mathbf{C}_T[T_1^*]\mathbf{C}_{T_1}[T^*] = \mathbf{C}_T[T^*] = I_{T^*}.$$

Taking determinants yields

$$\det(\mathbf{C}_T[T_1^*])\det(\mathbf{C}_{T_1}[T^*]) = 1.$$

Since all the matrices we have been considering have integer entries, the above determinants are integers. Hence

$$\det \mathbf{C}_T[T_1^*] = \pm 1,$$

as was to be proved.

A similar proof works for \mathbf{B}_T . □

11.3 Electrical networks

We will give a brief indication of the connection between the above discussion and the theory of electrical networks. Let D be a digraph, which for convenience we assume is *connected* and *loopless*. Suppose that at each edge e there is a voltage (potential difference) V_e from *init* e to *fin* e , and a current I_e in the direction of e (so a negative current I_e indicates a current of $|I_e|$ in the direction opposite to e). Think of V and I as functions on the edges, i.e., as elements of the vector space \mathbb{R}^E . There are three fundamental laws relating the quantities V_e and I_e .

Kirchhoff's First Law. $I \in \mathcal{C}_D$. In other words, the current flowing into a vertex equals the current flowing out. In symbols,

$$\sum_{\substack{e \\ \text{init } e=v}} I_e = \sum_{\substack{e \\ \text{fin } e=v}} I_e,$$


for all vertices $v \in V$.

Kirchhoff's Second Law. $V \in \mathcal{C}_D^\perp = \mathcal{B}$. In other words, the sum of the voltages around any circuit (called loops by electrical engineers), taking into account orientations, is 0.

Ohm's Law. *If edge e has resistance $R_e > 0$, then $V_e = I_e R_e$.*

The central problem of electrical network theory dealing with the above three laws² is the following: which of the $3q$ quantities V_e, I_e, R_e need to be specified to uniquely determine all the others, and how can we find or stipulate the solution in a fast and elegant way? We will be concerned here only with a special case, perhaps the most important special case in practical applications. Namely, suppose we apply a voltage V_q at edge e_q , with resistances R_1, \dots, R_{q-1} at the other edges e_1, \dots, e_{q-1} . Let V_i, I_i be the voltage and current at edge e_i . We would like to express each V_i and I_i in terms of V_q and R_1, \dots, R_{q-1} . By "physical intuition" there should be a unique solution, since we can actually build a network meeting the specifications of the problem. Note that if we have quantities V_i, I_i, R_i satisfying the three network laws above, then for any scalar α the quantities $\alpha V_i, \alpha I_i, R_i$ are also a solution. This means that we might as well assume that $V_q = 1$, since we can always multiply all voltages and currents afterwards by whatever value we want V_q to be.

As an illustration of a simple method of computing the total resistance of a network, the following diagram illustrates the notion of a *series connection* $D_1 + D_2$ and a *parallel connection* $D_1 \parallel D_2$ of two networks D_1 and D_2 with a distinguished edge e at which a voltage is applied.



When we apply a voltage $V_q > 0$ the current will flow along e_q from the higher potential to the lower. Thus $V_q/I_q < 0$, so we should define the *total resistance* $R(D)$ of the network D , together with the distinguished edge e , by $R(D) = -V_q/I_q$. It is well-known and easy to deduce from the three network


²Of course the situation becomes much more complicated when one introduces *dynamic* network elements like capacitors, alternating current, etc.

Laws that


$$R(D_1 + D_2) = R(D_1) + R(D_2)$$

$$\frac{1}{R(D_1 \parallel D_2)} = \frac{1}{R(D_1)} + \frac{1}{R(D_2)}.$$

A network that is built up from a single edge by a sequence of series and parallel connections is called a *series-parallel network*. An example is the following, with the distinguished edge e shown by a broken line from bottom to top.



The simplest network which is not a series-parallel network is called the *Wheatstone bridge* and is illustrated below. (The direction of the arrows has been chosen arbitrarily.) We will use this network as our main example in the discussion that follows.



We now return to an arbitrary connected loopless digraph D , with currents I_i , voltages V_i , and resistances R_i at the edges e_i . Recall that we are fixing $V_q = 1$ and R_1, \dots, R_{q-1} . Let T be a spanning tree of D . Since I is a

current if and only if it is orthogonal to the bond space \mathcal{B} (Theorem 11.3 and Kirchhoff's First Law), it follows that any basis for \mathcal{B} defines a complete and minimal set of linear relations satisfied by the I_i 's (namely, the relation that I is orthogonal to the basis elements). In particular, the basis matrix \mathbf{C}_T defines such a set of relations. For example, if D is the Wheatstone bridge shown above and if $T = \{e_1, e_2, e_5\}$, then we obtain the following relations by adding the edges e_1, e_2, e_5 of T in turn to T^* :

$$\begin{aligned} rclI_1 - I_3 - I_4 &= 0 \\ I_2 + I_3 + I_4 + I_6 &= 0 \\ I_4 + I_5 + I_6 &= 0. \end{aligned} \tag{11.5}$$

These three ($= p - 1$) equations give all the relations satisfied by the I_i 's alone, and the equations are linearly independent.

Similarly if V is a voltage then it is orthogonal to the cycle space \mathcal{C} . Thus any basis for \mathcal{C} defines a complete and minimal set of linear relations satisfied by the V_i 's (namely, the relation that V is orthogonal to the basis elements). In particular, the basis matrix \mathbf{C}_T defines such a set of relations. Continuing our example, we obtain the following relations by adding the edges e_3, e_4, e_6 of T^* in turn to T .

$$\begin{aligned} V_1 - V_2 + V_3 &= 0 \\ V_1 - V_2 + V_4 - V_5 &= 0 \\ V_2 + V_5 &= 1, \end{aligned} \tag{11.6}$$

These three ($= q - p + k$) equations give all the relations satisfied by the V_i 's alone, and the equations are linearly independent.

In addition, Ohm's Law gives the $q - 1$ equations $V_i = R_i I_i$, $1 \leq i \leq q - 1$. We have a total of $(p - k) + (q - p + k) + (q - 1) = 2q - 1$ equations in the $2q - 1$ unknowns I_i ($1 \leq i \leq q$) and V_i ($1 \leq i \leq q - 1$). Moreover, it is easy to see that these $2q - 1$ equations are linearly independent, using the fact that we already know that just the equations involving the I_i 's alone are linearly independent, and similarly the V_i 's. Hence this system of $2q - 1$ equations in $2q - 1$ unknowns has a unique solution. We have now reduced the problem to straightforward linear algebra. However, it is possible to describe the solution explicitly. We will be content here with giving a formula just for the total resistance $R(D) = -V_q/I_q = -1/I_q$.

Write the $2q - 1$ equations in the form of a $(2q - 1) \times 2q$ matrix \mathbf{K} . The columns of the matrix are indexed by $I_1, I_2, \dots, I_q, V_1, V_2, \dots, V_q$. The last

column V_q of the matrix keeps track of the constant terms of the equations. The rows of \mathbf{K} are given first by the equations among the I_i 's, then the V_i 's, and finally Ohm's Law. For our example of the Wheatstone bridge, we obtain the matrix

$$\mathbf{K} = \begin{array}{c|cccccc|cccccc} & I_1 & I_2 & I_3 & I_4 & I_5 & I_6 & V_1 & V_2 & V_3 & V_4 & V_5 & V_6 \\ \hline & 1 & 0 & -1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ & 0 & 1 & 1 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ & 0 & 0 & 0 & 0 & 0 & 0 & 1 & -1 & 1 & 0 & 0 & 0 \\ & 0 & 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 1 & -1 & 0 \\ & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & -1 & 1 \\ \hline R_1 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ & 0 & R_2 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ & 0 & 0 & R_3 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ & 0 & 0 & 0 & R_4 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 \\ & 0 & 0 & 0 & 0 & R_5 & 0 & 0 & 0 & 0 & 0 & -1 & 0 \end{array}$$

We want to solve for I_q by Cramer's rule. Call the submatrix consisting of all but the last column X . Let Y be the result of replacing the I_q column of X by the last column of \mathbf{K} . Cramer's rule then asserts that

$$I_q = \frac{\det Y}{\det X}.$$

We evaluate $\det X$ by taking a Laplace expansion along the first $p-1$ rows. In other words,

$$\det X = \sum_S \pm \det(X[[p-1], S]) \cdot \det(X[[p-1]^c, \bar{S}]), \quad (11.7)$$

where (a) S indexes all $(p-1)$ -element subsets of the columns, (b) $X[[p-1], S]$ denotes the submatrix of X consisting of entries in the first $p-1$ rows and in the columns S , (c) $X[[p-1]^c, \bar{S}]$ denotes the submatrix of X consisting of entries in the last $2q-p$ rows and in the columns other than S . In order for $\det(X[[p-1], S]) \neq 0$, we must choose $S = \{I_{i_1}, \dots, I_{i_{p-1}}\}$, where $\{e_{i_1}, \dots, e_{i_{p-1}}\}$ is a spanning tree T_1 (by Theorem 11.6(i)). In this case, $\det(X[[p-1], S]) = \pm 1$ by Theorem 11.13. If $I_q \notin S$, then the I_q column of $X[[p-1]^c, \bar{S}]$ will be zero. Hence to get a nonzero term in (11.7), we must

have $e_q \in S$. The matrix $X[[p-1]^c, \bar{S}]$ will have one nonzero entry in each of the first $q-p+1$ columns, namely, the resistances R_j where e_j is not an edge of T_1 . This accounts for $q-p+1$ entries from the last $q-1$ rows of $X[[p-1]^c, \bar{S}]$. The remaining $p-2$ of the last $q-1$ rows have available only one nonzero entry each, namely, a -1 in the columns indexed by V_j where e_j is an edge of T_1 other than e_q . Hence we need to choose $q-p+1$ remaining entries from rows p through q and columns indexed by V_j for e_j not an edge of T_1 . By Theorems 11.6(ii) and 11.13, this remaining submatrix has determinant ± 1 . It follows that

$$\det(X[[p-1], S]) \cdot \det(X[[p-1]^c, \bar{S}]) = \pm \prod_{e_j \notin E(T_1)} R_j.$$

Hence by (11.7), we get

$$\det X = \sum_{T_1} \pm \left(\prod_{e_j \notin E(T_1)} R_j \right), \quad (11.8)$$

where T_1 ranges over all spanning trees of D containing e_q . A careful analysis of the signs (omitted here) shows that all signs in (11.8) are negative, so we finally arrive at the remarkable formula

$$\det X = - \sum_{\substack{\text{spanning trees } T_1 \\ \text{containing } e_q}} \prod_{e_j \notin E(T_1)} R_j.$$

For example, if D is the Wheatstone bridge as above, and if we abbreviate $R_1 = a$, $R_2 = b$, $R_3 = c$, $R_4 = d$, $R_5 = e$, then

$$-\det X = abc + abd + abe + ace + ade + bcd + bde + cde.$$

Now suppose we replace column I_q in X by column V_q in the matrix \mathbf{K} , obtaining the matrix Y . There is a unique nonzero entry in the new column, so it must be chosen in any nonzero term in the expansion of $\det Y$. The argument now goes just as it did for $\det X$, except we have to choose S to correspond to a spanning tree T_1 that *doesn't* contain e_q . We therefore obtain

$$\det Y = \sum_{\substack{\text{spanning trees } T_1 \\ \text{not containing } e_q}} \prod_{\substack{e_j \notin E(T_1) \\ e_j \neq e_q}} R_j.$$

For example, for the Wheatstone bridge we get

$$\det Y = ac + ad + ae + bc + bd + be + cd + ce.$$

Recall that $I_q = \det(Y)/\det(X)$ and that the total resistance of the network is $-1/I_q$. Putting everything together gives our main result on electrical networks.

11.14 Theorem. *In the situation described above, the total resistance of the network is given by*

$$R(D) = -\frac{1}{I_q} = -\frac{\sum_{\substack{\text{spanning trees } T_1 \\ \text{containing } e_q}} \prod_{e_j \notin E(T_1)} R_j}{\sum_{\substack{\text{spanning trees } T_1 \\ \text{not containing } e_q}} \prod_{\substack{e_j \notin E(T_1) \\ e_j \neq e_q}} R_j}.$$

11.15 Corollary. *If the resistances R_1, \dots, R_{q-1} are all equal to one, then the total resistance of the network is given by*

$$R(D) = -\frac{1}{I_q} = \frac{\text{number of spanning trees containing } e_q}{\text{number of spanning trees not containing } e_q}.$$

In particular, if $R_1 = \dots = R_{q-1} = 1$, then the total resistance, when reduced to lowest terms a/b , has the curious property that the number $\kappa(D)$ of spanning trees of D is divisible by $a + b$.

11.4 Planar graphs (sketch)

A graph G is *planar* if it can be drawn in the plane \mathbb{R}^2 without crossing edges. A drawing of G in this way is called a *planar embedding*. An example of a planar embedding is shown in Figure 11.2. In this section we state the basic results on the bond and cycle spaces of a planar graph. The proofs are relatively straightforward and are omitted.

If the vertices and edges of a planar embedding of G are removed from \mathbb{R}^2 , then we obtain a disjoint union of open sets, called the *faces* (or *regions*) of G . (More precisely, these open sets are the faces of the planar embedding




Figure 11.2: A planar embedding

of G . Often we will not bother to distinguish between a planar graph and a planar embedding if no confusion should result.) Let $R = R(G)$ be the set of faces of G , and as usual $V(G)$ and $E(G)$ denote the set of vertices and edges of G , respectively.

NOTE. If G is a simple (no loops or multiple edges) planar embedding, then it can be shown that there exists a planar embedding of the same graph with edges as straight lines and with faces (regarding as the sequence of vertices and edges obtained by walking around the boundaries of the faces) preserved.

The *dual* G^* of the planar embedded graph G has vertex set $R(G)$ and edge set $E^*(G) = \{e^* : e \in E(G)\}$. If e is an edge of G , then let r and r' be the faces on its two sides. (Possibly $r = r'$; there are five such edges in Figure 11.2.) Then define e^* to connect r and r' . We can always draw G^* to be planar, letting e and e^* intersect once. If G is connected then every face of G^* contains exactly one (nonisolated) vertex of G and $G^{**} \cong G$. For any planar embedded graph G , the dual G^* is connected. Then $G \cong G^{**}$ if and only if G is connected. In general, we always have $G^* \cong G^{***}$. Figure 11.3 shows the dual G^* to the graph G of Figure 11.2, with the vertices of G^* drawn as open circles and the edges as broken lines.

11.16 Example. Let G consist of two disjoint edges. Then G^* has one vertex and two loops, while G^{**} is a three-vertex path. The unbounded face of G^* contains two vertices of G , and $G^{**} \not\cong G$.

Orient the edges of the planar graph G in any way to get a digraph D . Let r be an interior (i.e., bounded) face of D . An *outside edge* of r is an edge






Figure 11.3: A planar embedding and its dual

e such that r lies on one side of the edge, and a *different* face lies on the other side. The outside edges of any interior face r define a circulation (shown as solid edges in the diagram below), and these circulations (as r ranges over all interior faces of D) form a basis for the cycle space \mathcal{C}_G of G .



Given the orientation D of G , orient the edges of G^* as follows: as we walk along e in the direction of its orientation, e^* points to our *right*.



11.17 Theorem. Let $f: E(G) \rightarrow \mathbb{R}$. Define $f^*: E(G^*) \rightarrow \mathbb{R}$ by $f^*(e^*) = f(e)$. Then

$$\begin{aligned} f \in \mathcal{B}_G &\Leftrightarrow f^* \in \mathcal{C}_{G^*} \\ f \in \mathcal{C}_G &\Leftrightarrow f^* \in \mathcal{B}_{G^*}. \end{aligned}$$


11.18 Proposition. The set S is the set of edges of a spanning tree T of G if and only if $S^* = \{e^* : e \in S\}$ is the set of edges of a cotree T^* of G^* .

11.19 Corollary. $\kappa(G) = \kappa(G^*)$

For nonplanar graphs there is still a notion of a “dual” object, but it is no longer a graph but rather something called a *matroid*. Matroid theory is a flourishing subject which may be regarded as a combinatorial abstraction of linear algebra.

11.5 Squaring the square

A *squared rectangle* is a rectangle partitioned into finitely many (but more than one) squares. A squared rectangle is *perfect* if all the squares are of different sizes. The earliest perfect squared rectangle was found in 1936; its size is 33×32 and consists of nine squares:



The question then arose: does there exist a perfect squared square? A single example was found by Sprague in 1939; it has 55 squares. Then Brooks, Smith, Stone, and Tutte developed a network theory approach which we now explain.

The *Smith diagram* D of a squared rectangle is a directed graph whose vertices are the horizontal line segments of the squared rectangle and whose squares are the edges, directed from top to bottom. The top vertex (corresponding to the top edge of the rectangle being squared) and the bottom vertex (corresponding to the bottom edge) are called *poles*. Label each edge by the side length of the square to which it corresponds. Figure 11.4 shows the Smith diagram of the (perfect) squared rectangle above.

The following result concerning Smith diagrams is straightforward to verify.

11.20 Theorem. (a) If we set I_e and V_e equal to the label of edge e , then Kirchhoff's two laws hold (so $R_e = 1$) except at the poles.

(b) The Smith diagram is planar and can be drawn without separation of poles. Joining the poles by an edge from the bottom to the top gives a





Figure 11.4: A Smith diagram

3-connected graph, i.e., a connected graph that remains connected when one or two vertices are removed.

Call the 3-connected graph of Theorem 11.20 the *extended* Smith diagram of the $a \times b$ squared rectangle. If we impose a current $I_{e_1} = b$ on the new edge e_1 (directed from bottom to top) between poles, and a voltage $V_{e_1} = -a$, then Kirchhoff's two laws hold at *all* vertices. The diagram below shows the extended Smith diagram corresponding to Figure 11.4, with the new edge e_1 labelled by the current I_{e_1} .




We therefore have a recipe for searching for perfect squared rectangles and squares: start listing all 3-connected planar graphs. Then choose an edge e_1 to apply a voltage V_1 . Put a resistance $R_e = 1$ at the remaining edges e . Solve for $I_e (= V_e)$ to get a squared rectangle, and hope that one of these will be a square. One example Γ found by Brooks et al. was a 112×75 rectangle with 14 squares. It was given to Brooks' mother as a jigsaw puzzle,

and she found a different solution Δ ! We therefore have found a squared square (though not perfect):

Δ	75×75
112×112	Γ

Building on this idea, Brooks et al. finally found two 422×593 perfect rectangles with thirteen squares, all 26 squares being of different sizes. Putting them together as above gives a perfect squared square. This example has two defects: (a) it contains a smaller perfect squared rectangle (and is therefore not *simple*), and (b) it contains a “cross” (four squares meeting a point). They eventually found a perfect squared square with 69 squares without either of these defects. It is now known (thanks to computers) that the smallest order (number of squares) of a perfect squared square is 21. It is unique and happens to be simple and crossfree. See the figure below. It is known that the number (up to symmetry) of simple perfect squared squares of order n for $n \geq 21$ is 1, 8, 12, 26, 160, 441, 1152, . . .



Notes for Chapter 11

The theory of cycle spaces and bond spaces developed here had its origins with the pioneering work of G. Kirchhoff [64] in 1847. The proof given here of Theorem 11.13 is due to W. T. Tutte [119] in 1965. A nice account of the history of squaring the square due to Tutte appears in a *Scientific American* column by Martin Gardner [45]. See also [120] for another article by Tutte. A further survey article on this topic is by Kazarinoff and Weitzenkamp [63].

Chapter 12

Miscellaneous gems of algebraic combinatorics

12.1 The 100 prisoners

An evil warden is in charge of 100 prisoners (all with different names). He puts a row of 100 boxes in a room. Inside each box is the name of a different prisoner. The prisoners enter the room one at a time. Each prisoner must open 50 of the boxes, one at a time. If any of the prisoners does not see his or her own name, then they are all killed. The prisoners may have a discussion before the first prisoner enters the room with the boxes, but after that there is no further communication. A prisoner may not leave a message of any kind for another prisoner. In particular, all the boxes are shut once a prisoner leaves the room. If all the prisoners choose 50 boxes at random, then each has a success probability of $1/2$, so the probability that they are not killed is 2^{-100} , not such good odds. Is there a strategy that will increase the chances of success? What is the best strategy?

It's not hard to see that the prisoners can achieve a success probability of greater than 2^{-100} . For instance, suppose that the first prisoner opens the first 50 boxes and the second prisoner opens the last 50. If the first prisoner succeeds (with probability $1/2$), then the first prisoner's name is guaranteed not to be in one of the boxes opened by the second prisoner, so the second prisoner's probability of success is $50/99$. Each pair of prisoners can do this strategy, increasing the overall success probability to $(25/99)^{50}$, still an extremely low number. Can they do significantly better? The key to

understanding this problem is the realization that the prisoners do not have to decide in advance on which boxes they will open. A prisoner can decide which box to open next based on what he has seen in the boxes previously opened.

12.1 Theorem. *There exists a strategy with a success probability of*

$$1 - \sum_{j=51}^{100} \frac{1}{j} = 0.3118278207 \dots$$

Proof. The prisoners assign themselves the numbers $1, 2, \dots, 100$ by whatever method they prefer. Each prisoner is assigned a different number. The prisoners memorize everyone's number. They regard the boxes, which are lined up in a row, as being numbered $1, 2, \dots, 100$ from left-to-right. A prisoner with number k first goes to box k . If the prisoner sees his name, then he breathes a temporary sigh of relief, and the next prisoner enters. Otherwise the first prisoner will see the name of some other prisoner, say with number n_1 . He then opens box n_1 and repeats the procedure, so whenever he opens a box B that doesn't contain his own name, the next box that he opens has the number of the prisoner whose name appears in box B .

What is the probability of success of this strategy? Suppose that box i contains the name of the prisoner numbered $\pi(i)$. Thus π is a permutation of $1, 2, \dots, 100$. The boxes opened by prisoner i are those containing the names of prisoners with numbers $\pi(i), \pi^2(i), \pi^3(i)$, etc. If k is the length of the cycle of π containing i , then the prisoner will see his name after opening the k th box. This will happen whenever $k \leq 50$. Thus all prisoners see their names if and only if every cycle of π has length at most 50. If π does not have this property, then it has exactly one cycle of length $r > 50$. There are $\binom{100}{r}$ ways to choose the elements of the cycle and $(r-1)!$ ways to arrange them in a cycle. There are then $(100-r)!$ ways to arrange the other elements of π . Thus the number of permutations $\pi \in \mathfrak{S}_{100}$ with a cycle of length $r > 50$ is

$$\binom{100}{r} (r-1)! (100-r)! = \frac{100!}{r}.$$

(There are more clever ways to see this.) Hence the probability of success, i.e., the probability that π has *no* cycle of length more than 50, is

$$1 - \frac{1}{100!} \sum_{r=51}^{100} \frac{100!}{r} = 1 - \sum_{r=51}^{100} \frac{1}{r},$$

as claimed. \square

If we apply the above argument to $2n$ prisoners rather than 100, then we get a success probability of

$$1 - \sum_{r=n+1}^{2n} \frac{1}{r} = 1 - \sum_{r=1}^{2n} \frac{1}{r} + \sum_{r=1}^n \frac{1}{r}.$$

From calculus we know that there is a constant $\gamma = 0.577215665 \dots$, known as *Euler's constant*, for which

$$\lim_{n \rightarrow \infty} \left(\sum_{r=1}^n \frac{1}{r} - \log n \right) = \gamma.$$

It follows that as $n \rightarrow \infty$, the success probability of the prisoners is

$$\lim_{n \rightarrow \infty} (1 - \log 2n + \log n) = 1 - \log 2 = 0.3068528194 \dots.$$

It seems quite amazing that no matter how many prisoners there are, they can always achieve a success probability of over 30%!

NOTE. It can be shown that the above strategy is in fact *optimal*, i.e., no strategy achieves a higher probability of success. The proof, however, is not so easy.

12.2 Oddtown

The village of Oddtown has a population of n people. Inhabitants of Oddtown like to form clubs. Every club has an odd number of members, and every pair of clubs share an even number of members (possibly none).

12.2 Theorem. *There are at most n clubs.*

Proof. Let k be the number of clubs. Define a matrix $\mathbf{M} = (M_{ij})$ over the two-element field \mathbb{F}_2 as follows. The rows of \mathbf{M} are indexed by the clubs C_i and the columns by the inhabitants x_j of Oddtown. Set

$$M_{ij} = \begin{cases} 1, & x_j \in C_i \\ 0, & \text{otherwise.} \end{cases}$$

The matrix \mathbf{M} is called the *incidence matrix* corresponding to the clubs and their members.

In general, let S be a subset of $[n]$, and let $\chi_S \in \mathbb{Z}^n$ be the *characteristic vector* of S , i.e., $\chi_S = (a_1, \dots, a_n)$ where

$$a_i = \begin{cases} 1, & i \in S \\ 0, & i \notin S. \end{cases}$$

If T is another subset of $[n]$, then the key observation is that the scalar (dot) product of χ_S and χ_T is given by $\chi_S \cdot \chi_T = \#(S \cap T)$. Hence if we now work over \mathbb{F}_2 , then

$$\chi_S \cdot \chi_T = \begin{cases} 1, & \text{if } \#(S \cap T) \text{ is odd} \\ 0, & \text{if } \#(S \cap T) \text{ is even.} \end{cases} \quad (12.1)$$

Let $\mathbf{A} = \mathbf{M}\mathbf{M}^t$, a $k \times k$ matrix. By equation (12.1) and the assumption that every club has an odd number of members, we see that main diagonal elements of \mathbf{A} are 1. Similarly the off-diagonal elements of \mathbf{A} are 0, so $\mathbf{A} = I_k$, the $k \times k$ identity matrix. Hence $\text{rank}(\mathbf{A}) = k$.

Recall that if B is a $k \times m$ matrix and C is an $m \times n$ matrix (over some field), then $\text{rank}(BC) \leq \text{rank}(B)$ (as well as $\text{rank}(BC) \leq \text{rank}(C)$), since for any matrix D , $\text{rank}(D) = \dim \text{image}(D)$. Hence, since \mathbf{M} has n columns,

$$n \geq \text{rank}(\mathbf{M}) \geq \text{rank}(\mathbf{M}\mathbf{M}^t) = \text{rank}(\mathbf{A}) = k.$$

□

While Theorem 12.2 can be proved without linear algebra, the proof is not easy.

12.3 Complete bipartite partitions of K_n

Figure 12.1 shows the six edges of the complete graph K_4 partitioned (according to the edge label) into the edge sets of the three complete bipartite graphs $K_{3,1}$, $K_{2,1}$, and $K_{1,1}$. Clearly we can extend this construction, achieving a partition of the edges $E(K_n)$ of K_n into the edge sets of $n - 1$ complete bipartite graphs. Specifically, let E_1 be the set of edges incident to a fixed vertex v . Thus E_1 is the edge set of a complete bipartite graph $K_{n-1,1}$. Remove E_1 from $E(K_n)$ and proceed by induction, obtaining a partition of $E(K_n)$ into the edges of $K_{n-1,1}$, $K_{n-2,1}$, \dots , $K_{1,1}$. The question thus arises




Figure 12.1: A decomposition of the edges of K_4 into three complete bipartite graphs

as to whether $E(K_n)$ can be partitioned into *fewer* than $n - 1$ edge sets of complete bipartite graphs.

12.3 Theorem. *If $E(K_n)$ is the disjoint union of the edge sets of m complete bipartite graphs, then $m \geq n - 1$.*

Proof. Let $E(K_n) = E(B_1) \cup E(B_2) \cup \cdots \cup E(B_m)$ (disjoint union), where B_k is a complete bipartite graph with vertex bipartition (X_k, Y_k) (so $X_k \cap Y_k = \emptyset$). For $1 \leq i \leq n$, define an $n \times n$ matrix A_k by

$$(A_k)_{ij} = \begin{cases} 1, & i \in X_k, j \in Y_k \\ 0, & \text{otherwise.} \end{cases}$$

All nonzero rows of A_k are equal, so $\text{rank } A_k = 1$. Let $S = \sum_{k=1}^m A_k$. For $i \neq j$, exactly one of the $2m$ numbers $(A_k)_{ij}$ and $(A_k)_{ji}$, $1 \leq k \leq m$, is equal to 1, since every edge ij of K_n appears in one $E(B_k)$ with either $i \in X_k$ and $j \in Y_k$, or else $j \in X_k$ and $i \in Y_k$. Hence

$$S + S^t = J - I,$$

where as usual J is the $n \times n$ all 1's matrix, and I is the $n \times n$ identity matrix.

Claim. If T is any real matrix satisfying $T + T^t = J - I$, then $\text{rank } T \geq n - 1$.

Suppose to the contrary that $\text{rank } T \leq n - 2$. Then T has (at least) two linearly independent eigenvectors x, y such that $Tx = Ty = 0$ [why?]. Since J has rank one, the space $\langle x, y \rangle$ spanned by x and y contains a nonzero vector z satisfying $Jz = 0$ [why?]. Then from $T + T^t = J - I$ and $Tz = 0$ we get $-z = T^t z$. Take the dot product with z^t on the left. We get

$$\begin{aligned} -|z|^2 &= z^t T^t z \\ &= (z^t T^t z)^t \quad (\text{since a } 1 \times 1 \text{ matrix is symmetric}) \\ &= z^t T z \quad (\text{since in general } (AB)^t = B^t A^t) \\ &= 0 \quad (\text{since } Tz = 0), \end{aligned}$$

contradicting $z \neq 0$. Hence the claim is proved, so in particular $\text{rank } X \geq n - 1$. But in general $\text{rank}(A + B) \leq \text{rank } A + \text{rank } B$ [why?]. Therefore from $S = \sum_{k=1}^m A_k$ and $\text{rank } A_k = 1$ we get $\text{rank } S \leq m$. It follows that $m \geq n - 1$, completing the proof. \square

12.4 The nonuniform Fisher inequality

A *balanced incomplete block design* (BIBD) with parameters (v, k, λ, r, b) is a v -element set X and a collection \mathcal{A} of k -element subsets (blocks), with $\#\mathcal{A} = b$, such that any two points $x, y \in X$ lie in exactly λ blocks, and each point is in exactly r blocks. We also assume that $k < v$, which is the reason for the word “incomplete.” We can draw a BIBD as a bipartite graph with vertex bipartition (X, \mathcal{A}) . There is an edge from $x \in X$ to $A \in \mathcal{A}$ if $x \in A$. Thus the degree of each vertex $x \in X$ is r , and the degree of each vertex $A \in \mathcal{A}$ is k . It follows that $vr = kb$ (the total number of edges of the graph). We can also count the number of two-element sets of edges that are incident to the same vertex of \mathcal{A} . On the one hand, since each vertex in \mathcal{A} has degree k this number is $b \binom{k}{2}$. On the other hand, each pair of points in X are mutually adjacent to λ points in \mathcal{A} , so we get $\lambda \binom{v}{2} = b \binom{k}{2}$. A little manipulation shows that the two equalities $vr = kb$ and $\lambda \binom{v}{2} = b \binom{k}{2}$ are equivalent to

$$vr = kb, \quad \lambda(v - 1) = r(k - 1),$$

the usual form in which they are written.

R. A. Fisher showed in 1940 that $b \geq v$. This inequality was generalized by R. C. Bose in 1949. The most convenient way to state Bose’s inequalities,

known as the *nonuniform Fisher inequality*, is to reverse the roles of points and blocks. Thus consider the elements x of X to be sets whose elements are the blocks $A \in \mathcal{A}$ that contain them. In other words, we have a collection C_1, \dots, C_v of r -element sets whose union contains b points x_1, \dots, x_b . Each point is in exactly k of the sets. Finally, $\#(C_i \cap C_j) = \lambda$ for all $i \neq j$.

12.4 Theorem. *Let C_1, \dots, C_v be distinct subsets of a b -element set X such that for all $i \neq j$ we have $\#(C_i \cap C_j) = \lambda$ for some $1 \leq \lambda < b$ (independent of i and j). Then $v \leq b$.*

Proof. Case 1: some $\#C_i = \lambda$. Then all other C_j 's contain C_i and are disjoint otherwise, so

$$v \leq \underbrace{1}_{\text{from } C_i} + \underbrace{b - \lambda}_{\text{from all } C_j \neq C_i} \leq b.$$

Case 2: all $\#C_i > \lambda$. Let $\gamma_i = \#C_i - \lambda > 0$. Let \mathbf{M} be the incidence matrix of the set system C_1, \dots, C_v , i.e., the rows of \mathbf{M} correspond to the C_i 's and the columns to the elements x_1, \dots, x_b of X , with

$$M_{ij} = \begin{cases} 1, & x_j \in C_i \\ 0, & x_j \notin C_i. \end{cases}$$

Let $\mathbf{A} = \mathbf{M}\mathbf{M}^t$. The hypotheses imply that $\mathbf{A} = \lambda\mathbf{J} + \mathbf{G}$, where \mathbf{J} as usual is the all 1's matrix (of size v), and \mathbf{G} is the diagonal matrix $\text{diag}(\gamma_1, \dots, \gamma_v)$.

Claim: $\text{rank}(\mathbf{A}) = v$ (i.e., \mathbf{A} is invertible). We would then have

$$v = \text{rank}(\mathbf{A}) \leq \text{rank}(\mathbf{M}) \leq b,$$


the last inequality because \mathbf{M} has b columns.

As in the proof of Theorem 4.7, a real symmetric matrix \mathbf{B} is positive semidefinite if it has nonnegative eigenvalues. Equivalently, by basic linear algebra, $u\mathbf{B}u^t \geq 0$ for all row vectors u of length v . Moreover \mathbf{B} is positive definite (and so has positive eigenvalues) if $u\mathbf{B}u^t > 0$ for all $u \neq 0$.

Now we easily compute that

$$u(\lambda\mathbf{J} + \mathbf{G})u^t = \lambda(u_1 + \dots + u_v)^2 + \gamma_1 u_1^2 + \dots + \gamma_v u_v^2 > 0$$

for all $u \neq 0$. Thus $\mathbf{A} = \lambda\mathbf{J} + \mathbf{G}$ is positive definite and hence of full rank v . \square

Figure 12.2: The 3×4 grid graph

12.5 Odd neighborhood covers

Consider an $m \times n$ grid graph. The case $m = 3$, $n = 4$ is shown in Figure 12.2. At each vertex is a turned on light bulb and also a switch that changes the state of its bulb and those of its neighbors (adjacent vertices). Can all the lights be turned off?

This problem was open for many years until in 1989 K. Sutner, then a graduate student, showed using automata theory that the answer is yes for *any* (finite) graph! More explicitly, let G be a finite graph with a turned on light bulb at each vertex. At each vertex is a switch that changes the state of that vertex and all its neighbors. Then it is possible to turn off all the lights. We will give a modification of a simpler proof due to Y. Caro based on linear algebra.

Without loss of generality we may assume that G is simple. If $v \in V(G)$, then the *neighborhood* $N(v)$ of v is the set consisting of v and all vertices adjacent to v . A little thought shows that we need to prove the following result.

12.5 Theorem. *There exists a subset $S \subseteq V = V(G)$ such that $\#(S \cap N(v))$ is odd for all $v \in V$. (It follows that switching at the vertices $v \in S$ turns all the lights off.)*

Proof. Let $V(G) = \{v_1, \dots, v_p\}$. Let \mathbf{A} be the adjacency matrix of G over the field \mathbb{F}_2 , and let $y = (1, 1, \dots, 1) \in \mathbb{F}_2^p$. Write $\text{row}(\mathbf{B})$ for the row space of a matrix \mathbf{B} . Given $S \subseteq V$, let $\chi_S = (a_1, \dots, a_p) \in \mathbb{F}_2^p$ denote the characteristic (row) vector of S , i.e.,

$$a_i = \begin{cases} 1, & v_i \in S \\ 0, & v_i \notin S. \end{cases}$$

Note that switching at S turns all the lights off if and only if $\chi_S(\mathbf{A} + I) = y$. Hence we need to show that $y \in \text{row}(\mathbf{A} + I)$ [why?].

Let us recall from linear algebra some standard facts about orthogonal subspaces. Let K be a field, and for $u, v \in K^n$ let $u \cdot v$ be the usual dot product (2.1) of u and v , so $u \cdot v \in K$. If W is a subspace of K^n , then define the *orthogonal subspace* W^\perp by

$$W^\perp = \{u \in K^n : u \cdot v = 0 \text{ for all } v \in W\}.$$

(In Chapter 11 we discussed the case $K = \mathbb{R}$.) Let $d = \dim W$. Since W^\perp is the set of solutions to d linearly independent homogeneous linear equations [why?], we have

$$\dim W + \dim W^\perp = n. \quad (12.2)$$

Note that by definition of $^\perp$ we have $W \subseteq (W^\perp)^\perp$. By equation (12.2) and the equation obtained from it by replacing W with W^\perp , we get $\dim W = \dim (W^\perp)^\perp$. Hence

$$(W^\perp)^\perp = W. \quad (12.3)$$

NOTE. Though irrelevant here, let us point out that if $K \subseteq \mathbb{R}$ then $W \cap W^\perp = \{0\}$, but that this fact need not hold in characteristic $p \neq 0$. Over \mathbb{C} we should define $u \cdot v = u_1 \bar{v}_1 + \cdots + u_n \bar{v}_n$, where $\bar{}$ denotes complex conjugation, in order to get the most sensible theory.

Now by equation (12.3) the vector $y = (1, 1, \dots, 1)$ (or any vector in \mathbb{F}_2^n) lies in the row space of $\mathbf{A} + I$ if and only if it is orthogonal to every vector in $\text{row}(\mathbf{A} + I)^\perp = \ker(\mathbf{A} + I)$. Thus we need to show that if $(\mathbf{A} + I)v^t = 0$, then $v \cdot y = 0$. Equivalently, if $yv^t \neq 0$ then $(\mathbf{A} + I)v^t \neq 0$. Note that (a) $yv^t \neq 0$ means that v has an odd number of 1's, and (b) $(\mathbf{A} + I)v^t$ is the sum of the rows of $\mathbf{A} + I$ indexed by the positions of the 1's in v . Thus we need to show that $\mathbf{A} + I$ does not have an odd number of rows summing to 0.

Suppose that v_1, \dots, v_k are vertices indexing rows of \mathbf{A} summing to 0. Let H be the subgraph induced by v_1, \dots, v_k , i.e., H consists of the vertices v_1, \dots, v_k and all edges of G between two of these vertices. Let b_{ij} be the (i, j) -entry of $\mathbf{A} + I$. Since $\sum_{i=1}^k b_{ij} = 0$ for $1 \leq j \leq n$, and each $b_{ii} = 1$, it follows that every vertex of H has odd degree. Since [why?]

$$\sum_{v \in V(H)} \deg v = 2 \cdot \#E(H),$$

we have that $k = \#V(H)$ is even, completing the proof. \square

12.6 Circulant Hadamard matrices

For our next “gem of algebraic combinatorics,” we will provide some variety by leaving the realm of linear algebra and looking at some simple algebraic number theory.

An $n \times n$ matrix H is a *Hadamard matrix* if its entries are ± 1 and its rows are orthogonal. Equivalently, its entries are ± 1 and $HH^t = nI$. In particular [why?],

$$\det H = \pm n^{n/2}. \quad (12.4)$$

It is easy to see that if H is an $n \times n$ Hadamard matrix then $n = 1$, $n = 2$, or $n = 4m$ for some integer $m \geq 1$. (See Exercise 12.18.) It is conjectured that the converse is true, i.e., for every such n there exists an $n \times n$ Hadamard matrix.

An $n \times n$ matrix $A = (b_{ij})$ is a *circulant* or *circulant matrix* if it has the form $b_{ij} = a_{i-j}$ for some a_0, a_1, \dots, a_{n-1} , where the subscript $i - j$ is taken modulo n . For instance,

$$A = \begin{bmatrix} a & b & c & d \\ d & a & b & c \\ c & d & a & b \\ b & c & d & a \end{bmatrix}$$

is a circulant. Let $A = (a_{i-j})$ be an $n \times n$ circulant, and let $\zeta = e^{2\pi i/n}$, a primitive n th root of unity. It is straightforward to compute that for $0 \leq j < n$ the column vector $[1, \zeta^j, \zeta^{2j}, \dots, \zeta^{(n-1)j}]^t$ is an eigenvector of A with eigenvalue $a_0 + \zeta^j a_1 + \zeta^{2j} a_2 + \dots + \zeta^{(n-1)j} a_{n-1}$. Hence

$$\det A = \prod_{j=0}^{n-1} (a_0 + \zeta^j a_1 + \zeta^{2j} a_2 + \dots + \zeta^{(n-1)j} a_{n-1}). \quad (12.5)$$

Note that the matrix

$$\begin{bmatrix} -1 & 1 & 1 & 1 \\ 1 & -1 & 1 & 1 \\ 1 & 1 & -1 & 1 \\ 1 & 1 & 1 & -1 \end{bmatrix}$$

is both a Hadamard matrix and a circulant.

Conjecture. Let H be an $n \times n$ circulant Hadamard matrix. Then $n = 1$ or $n = 4$.

The first significant work on this conjecture is due to R. J. Turyn. He showed that there does not exist a circulant Hadamard matrix of order $8m$, and he also excluded certain other orders of the form $4(2m + 1)$. Turyn's proofs use the machinery of algebraic number theory. Here we will give a proof for the special case $n = 2^k$, $k \geq 3$, where the algebraic number theory can be "dumbed down" to elementary commutative algebra and field theory. (Only in Theorem 12.14 do we use a little Galois theory, which can be avoided with a bit more work.) It would be interesting to find similar proofs for other values of n .

12.6 Theorem. *There does not exist a circulant Hadamard matrix H of order 2^k , $k \geq 3$.*

NOTE. It is curious that the numbers 2^k ($k \geq 2$) are the easiest multiples of 4 to show are *not* the orders of circulant Hadamard matrices, while on the other hand the numbers 2^k ($k \geq 1$) are the easiest numbers to show *are* the orders of Hadamard matrices. To see that 2^k is the order of a Hadamard matrix H , first note that the case $k = 1$ is trivial. It is routine to show that if H_1 is a Hadamard matrix of order a and H_2 is a Hadamard matrix of order b , then the tensor (or Kronecker) product $A \otimes B$ is a Hadamard matrix of order ab . It follows that there exists a Hadamard matrix of order 2^k , $k \geq 1$.

From now on we assume $n = 2^k$ and $\zeta = e^{2\pi i/2^k}$. Clearly ζ is a zero of the polynomial $p_k(x) = x^{2^{k-1}} + 1$. We will be working in the ring $\mathbb{Z}[\zeta]$, the smallest subring of \mathbb{C} containing \mathbb{Z} and ζ . Write $\mathbb{Q}(\zeta)$ for the quotient field of $\mathbb{Z}[\zeta]$, i.e., the field obtained by adjoining ζ to \mathbb{Q} .

12.7 Lemma. *The polynomial $p_k(x)$ is irreducible over \mathbb{Q} .*

Proof. If $p_k(x)$ is reducible then so is $p_k(x + 1)$. A standard fact about polynomial factorization is *Gauss' lemma*, namely, an integral polynomial that factors over \mathbb{Q} also factors over \mathbb{Z} . If $p(x), q(x) \in \mathbb{Z}[x]$, write $p(x) \equiv q(x) \pmod{2}$ to mean that the coefficients of $p(x) - q(x)$ are even. Now [why?]

$$p_k(x + 1) \equiv (x + 1)^{2^{k-1}} + 1 \equiv x^{2^{k-1}} \pmod{2}.$$

Hence any factorization of $p_k(x + 1)$ over \mathbb{Z} into two factors of degree at least one has the form $p_k(x + 1) = (x^r + 2a)(x^s + 2b)$, where $r + s = 2^{k-1}$ and a, b

are polynomial of degrees less than r and s , respectively. Hence the constant term of $p_k(x+1)$ is divisible by 4, a contradiction. \square

It follows by elementary field theory that every element $u \in \mathbb{Z}[\zeta]$ can be uniquely written in the form

$$u = b_0 + b_1\zeta + b_2\zeta^2 + \cdots + b_{n/2-1}\zeta^{n/2-1}, \quad b_i \in \mathbb{Z}.$$

The basis for our proof of Theorem 12.6 is the two different ways to compute $\det H$ given by equations (12.4) and (12.5), yielding the formula

$$\prod_{j=0}^{n-1} (a_0 + \zeta^j a_1 + \zeta^{2j} a_2 + \cdots + \zeta^{(n-1)j} a_{n-1}) = \pm n^{n/2} = \pm 2^{k2^{k-1}}. \quad (12.6)$$

Thus we have a factorization in $\mathbb{Z}[\zeta]$ of $2^{k2^{k-1}}$. Algebraic number theory is concerned with factorization of algebraic integers (and ideals) in algebraic number fields, so we have a vast amount of machinery available to show that no factorization (12.6) is possible (under the assumption that each $a_j = \pm 1$). Compare Kummer's famous approach toward Fermat's Last Theorem (which led to his creation of algebraic number theory), in which he considered the equation $x^n + y^n = z^n$ as $\prod_{\tau^n=1} (x + \tau y) = z^n$.

We are continuing to assume that $H = (a_{j-i})$ is an $n \times n$ circulant Hadamard matrix. We will denote the eigenvalues of H by

$$\gamma_j = a_0 + a_1\zeta^j + a_2\zeta^{2j} + \cdots + a_{n-1}\zeta^{(n-1)j}.$$

12.8 Lemma. *For $0 \leq j \leq n-1$ we have*

$$|\gamma_j| = \sqrt{n}.$$

Thus all the factors appearing on the left-hand side of (12.6) have absolute value \sqrt{n} .

First proof (naive). Let H_i denote the i th row of H , let \cdot denote the usual dot product, and let $\bar{\cdot}$ denote complex conjugation. Then

$$\begin{aligned} \gamma_j \bar{\gamma}_j &= (a_0 + a_1\zeta^j + \cdots + a_{n-1}\zeta^{(n-1)j})(a_0 + a_1\zeta^{-j} + \cdots + a_{n-1}\zeta^{-(n-1)j}) \\ &= H_1 \cdot H_1 + (H_1 \cdot H_2)\zeta^j + (H_1 \cdot H_3)\zeta^{2j} + \cdots + (H_1 \cdot H_n)\zeta^{(n-1)j}. \end{aligned}$$

By the Hadamard property we have $H_1 \cdot H_1 = n$, while $H_1 \cdot H_k = 0$ for $2 \leq k \leq n$, and the proof follows. \square

Second proof (algebraic). The matrix $\frac{1}{\sqrt{n}}H$ is a real orthogonal matrix. By linear algebra, all its eigenvalues have absolute value 1. Hence all eigenvalues γ_j of H have absolute value \sqrt{n} . \square

12.9 Lemma. *We have*

$$2 = (1 - \zeta)^{n/2}u, \quad (12.7)$$

where u is a unit in $\mathbb{Z}[\zeta]$.

Proof. Put $x = 1$ in

$$x^{n/2} + 1 = \prod_{\substack{j=0 \\ j \text{ odd}}}^{n-1} (x - \zeta^j)$$

to get $2 = \prod_j (1 - \zeta^j)$. Since

$$1 - \zeta^j = (1 - \zeta)(1 + \zeta + \cdots + \zeta^{j-1}),$$

it suffices to show that $1 + \zeta + \cdots + \zeta^{j-1}$ is a unit when j is odd. Let $j\bar{j} \equiv 1 \pmod{n}$. Then

$$\begin{aligned} (1 + \zeta + \cdots + \zeta^{j-1})^{-1} &= \frac{1 - \zeta}{1 - \zeta^j} \\ &= \frac{1 - (\zeta^j)^{\bar{j}}}{1 - \zeta^j} \in \mathbb{Z}[\zeta], \end{aligned}$$

as desired. \square

12.10 Lemma. *We have $\mathbb{Z}[\zeta]/(1 - \zeta) \cong \mathbb{F}_2$.*

Proof. Let $R = \mathbb{Z}[\zeta]/(1 - \zeta)$. The integer 2 is not a unit in $\mathbb{Z}[\zeta]$, e.g., because $1/2$ is not an algebraic integer (the zero of a *monic* polynomial $f(x) \in \mathbb{Z}[x]$). Thus by Lemma 12.9, $1 - \zeta$ is also not a unit. Hence $R \neq 0$.

For all j we have $\zeta^j = 1$ in R since $\zeta = 1$ in R . Hence all elements of R can be written as ordinary integers m . But $0 = 2$ in R by Lemma 12.9, so the only elements of R are 0 and 1. \square

12.11 Lemma. *For all $0 \leq j \leq n - 1$ there is an integer $h_j \geq 0$ such that*

$$a_0 + a_1\zeta^j + a_2\zeta^{2j} + \cdots + a_{n-1}\zeta^{(n-1)j} = v_j(1 - \zeta)^{h_j},$$

where v_j is a unit in $\mathbb{Z}[\zeta]$.

Proof. Since 2 is a multiple of $1 - \zeta$ by Lemma 12.9, we have by (12.6) that

$$\prod_{j=0}^{n-1} (a_0 + a_1\zeta^j + a_2\zeta^{2j} + \cdots + a_{n-1}\zeta^{(n-1)j}) = 0$$

in $\mathbb{Z}[\zeta]/(1 - \zeta)$. Since $\mathbb{Z}[\zeta]/(1 - \zeta)$ is an integral domain by Lemma 12.10, some factor $a_0 + a_1\zeta^j + \cdots + a_{n-1}\zeta^{(n-1)j}$ is divisible by $1 - \zeta$. Divide this factor and the right-hand side of (12.6) by $1 - \zeta$, and iterate the procedure. We continue to divide a factor of the left-hand side and the right-hand side by $1 - \zeta$ until the right-hand side becomes the unit u . Hence each factor of the original product has the form $v(1 - \zeta)^h$, where v is a unit. \square

12.12 Corollary. *Either $\gamma_0/\gamma_1 \in \mathbb{Z}[\zeta]$ or $\gamma_1/\gamma_0 \in \mathbb{Z}[\zeta]$. (In fact, both $\gamma_0/\gamma_1 \in \mathbb{Z}[\zeta]$ and $\gamma_1/\gamma_0 \in \mathbb{Z}[\zeta]$, as will soon become apparent, but we don't need this fact here.)*

Proof. By the previous lemma, each γ_j has the form $v_j(1 - \zeta)^{h_j}$. If $h_0 \geq h_1$ then $\gamma_0/\gamma_1 \in \mathbb{Z}[\zeta]$; otherwise $\gamma_1/\gamma_0 \in \mathbb{Z}[\zeta]$. \square

We now need to appeal to a result of Kronecker on elements of $\mathbb{Z}[\zeta]$ of absolute value one. For completeness we include a proof of this result, beginning with a lemma. Recall that if θ is an algebraic number (the zero of an irreducible polynomial $f(x) \in \mathbb{Q}[x]$), then a *conjugate* of θ is any zero of $f(x)$.

12.13 Lemma. *Let θ be an algebraic integer such that θ and all its conjugates have absolute value one. Then θ is a root of unity.*

Proof. Suppose the contrary. Let $\deg \theta = d$, i.e., $[\mathbb{Q}(\theta) : \mathbb{Q}] := \dim_{\mathbb{Q}} \mathbb{Q}(\theta) = d$. Now $\theta, \theta^2, \theta^3, \dots$ are all distinct and hence infinitely many of them have the property that no two are conjugate. Each $\theta^j \in \mathbb{Z}[\theta]$ and so is the root of a monic integral polynomial of degree at most d , since the set of algebraic integers forms a ring. If $\theta_1, \theta_2, \dots, \theta_d$ are the conjugates of θ , then all the conjugates of θ^j are among $\theta_1^j, \theta_2^j, \dots, \theta_d^j$. Hence each θ^j satisfies the hypothesis that all its conjugates have absolute value 1 (and θ^j is an algebraic integer). Thus the r th elementary symmetric function e_r in θ^j and its conjugates has at most $\binom{d}{r}$ terms, each of absolute value 1, so $|e_r| \leq \binom{d}{r}$. Moreover, $e_r \in \mathbb{Z}$ since θ^j is an algebraic integer. It follows that there are only finitely many possible polynomials that can be an irreducible monic polynomial with a zero equal to some θ^j , contradicting the fact that there are infinitely many θ^j 's for which no two are conjugate. \square

12.14 Theorem (Kronecker). *Let τ be any root of unity and $\alpha \in \mathbb{Z}[\tau]$ with $|\alpha| = 1$. Then α is a root of unity.*

Proof. Since $\alpha \in \mathbb{Z}[\tau]$, we see that α is an algebraic integer. We use the basic fact from Galois theory that the Galois group of the extension field $\mathbb{Q}(\tau)/\mathbb{Q}$ is abelian. Let β be a conjugate of α , so $\beta = w(\alpha)$ for some automorphism w of $\mathbb{Q}(\tau)$. Apply w to the equation $\alpha\bar{\alpha} = 1$. Since complex conjugation is an automorphism of $\mathbb{Q}(\tau)$ it commutes with w , so we obtain $\beta\bar{\beta} = 1$. Hence all the conjugates of α have absolute value one, so α is a root of unity by the previous lemma. \square

For our next result, we need the standard algebraic fact that if $\tau = e^{2\pi i/m}$, a primitive m th root of unity, then $[\mathbb{Q}(\tau) : \mathbb{Q}] = \phi(m)$ (the Euler ϕ -function). Equivalently, the unique monic polynomial $\Phi_m(x)$ whose zeros are the primitive m th roots of unity is irreducible. This polynomial is by definition given by

$$\Phi_m(x) = \sum_{\substack{1 \leq j \leq m \\ \gcd(j,m)=1}} (x - \tau^j)$$

and is called a *cyclotomic polynomial*. Lemma 12.7 is the case $m = n (= 2^k)$.

12.15 Lemma. *If $\tau \in \mathbb{Z}[\zeta]$ is a root of unity, then $\tau = \zeta^r$ for some $r \in \mathbb{Z}$.*

Proof. Suppose not. It is easy to see that then either τ is a primitive 2^m th root of unity for some $m > k$, or else τ^s is a primitive p th root of unity for some odd prime p and some $s \geq 1$. In the former case

$$[\mathbb{Q}(\tau) : \mathbb{Q}] = \phi(2^m) = 2^{m-1} > 2^{k-1} = \phi(2^k) = [\mathbb{Q}(\zeta) : \mathbb{Q}],$$

a contradiction. In the latter case, $\tau^s \zeta$ is a primitive pn -th root of unity, so

$$[\mathbb{Q}(\tau^s \zeta) : \mathbb{Q}] = \phi(pn) = (p-1)\phi(n) > \phi(n) = [\mathbb{Q}(\zeta) : \mathbb{Q}],$$

again a contradiction. \square

We now have all the ingredients to complete the proof of Theorem 12.6. Note that we have yet to use the hypothesis that $a_i = \pm 1$. By Lemma 12.8 we have

$$|\gamma_1/\gamma_0| = |\gamma_0/\gamma_1| = 1.$$

Hence by Corollary 12.12, Theorem 12.14 and Lemma 12.15 we have $\gamma_0 = \zeta^{-r}\gamma_1$ for some r . Expand γ_0 and $\zeta^{-r}\gamma_1$ uniquely as integer linear combinations of $1, \zeta, \zeta^2, \dots, \zeta^{\frac{n}{2}-1}$:

$$\begin{aligned}\gamma_0 &= a_0 + a_1 + \cdots + a_{n-1} = \pm\sqrt{n} \\ \zeta^{-r}\gamma_1 &= \zeta^{-r}((a_0 - a_{n/2}) + (a_1 - a_{n/2+1})\zeta + \cdots) \\ &= (a_r - a_{n/2+r}) + (a_{r+1} - a_{n/2+r+1})\zeta + \cdots.\end{aligned}$$

Equating coefficients of ζ^0 yields $\pm\sqrt{n} = a_r - a_{n/2+r}$. Since each $a_i = \pm 1$, we must have $n \leq 4$, completing the proof. \square

12.7 P -recursive functions

A function $f: \mathbb{N} \rightarrow \mathbb{C}$ is called *polynomially recursive*, or *P -recursive* for short, if there exist polynomials $P_0(n), \dots, P_d(n) \in \mathbb{C}[n]$, with $P_d(n) \neq 0$, such that

$$P_d(n)f(n+d) + P_{d-1}(n)f(n+d-1) + \cdots + P_0(n)f(n) = 0 \quad (12.8)$$

for all $n \geq 0$.

For instance, the Fibonacci sequence F_n is P -recursive since $F_{n+2} - F_{n+1} - F_n = 0$ for all $n \geq 0$. Here $d = 2$ and $P_2(n) = 1$, $P_1(n) = P_0(n) = -1$. This situation is quite special since the polynomials $P_i(n)$ are *constants*. Another P -recursive function is $f(n) = n!$, since $f(n+1) - (n+1)f(n) = 0$ for all $n \geq 0$.

Let \mathcal{P} denote the set of all P -recursive functions $f: \mathbb{N} \rightarrow \mathbb{C}$. Our goal in this section is to prove that \mathcal{P} is a \mathbb{C} -algebra, that is, for any $f, g \in \mathcal{P}$ and $\alpha, \beta \in \mathbb{C}$, we have

$$\alpha f + \beta g \in \mathcal{P}, \quad fg \in \mathcal{P},$$

with obvious compatibility properties such as $(\alpha f)g = f(\alpha g) = \alpha(fg)$. There is one technical problem that needs to be dealt with before proceeding to the proof. We would like to conclude from equation (12.8) that

$$f(n+d) = -\frac{1}{P_d(n)}(P_{d-1}(n)f(n+d-1) + \cdots + P_0(n)f(n)). \quad (12.9)$$

This formula, however, is problematical when $P_d(n) = 0$. This can happen only for finitely many n , so equation (12.9) is valid for n sufficiently large.

Thus we want to deal with functions $f(n)$ only for n sufficiently large. To this end, define $f \sim g$ if $f(n) = g(n)$ for all but finitely many n . Clearly \sim is an equivalence relation; the equivalence classes are called *germs* at ∞ of functions $f: \mathbb{N} \rightarrow \mathbb{C}$. The germ containing f is denoted $[f]$. Write \mathcal{G} for the set of all germs.

12.16 Lemma. (a) *If f is P -recursive and $f \sim g$, then g is P -recursive. In other words, the property of P -recursiveness is compatible with the equivalence relation \sim .*

(b) *Write $\mathbb{C}^{\mathbb{N}}$ for the complex vector space of all functions $f: \mathbb{N} \rightarrow \mathbb{C}$. Let $\alpha, \beta \in \mathbb{C}$ and $f_1, f_2, g_1, g_2 \in \mathbb{C}^{\mathbb{N}}$. If $f_1 \sim f_2$ and $g_1 \sim g_2$, then $\alpha f_1 + \beta g_1 \sim \alpha f_2 + \beta g_2$ and $f_1 g_1 \sim f_2 g_2$. In other words, linear combinations and multiplication are compatible with the equivalence relation \sim . Thus the set \mathcal{G} has the structure of a \mathbb{C} -algebra, i.e., a complex vector space and a ring (with obvious compatibility properties such as $(\alpha f)g = f(\alpha g) = \alpha(fg)$).*

Proof. (a) Suppose that $f(n) = g(n)$ for all $n > n_0$. Let (12.8) be the recurrence satisfied by f . Multiply both sides by $\prod_{j=0}^{n_0} (n-j)$. We then get a recurrence relation satisfied by g . Hence g is P -recursive.

(b) This is clear. □

Let $\mathbb{C}[n]$ denote the ring of complex polynomials in n . Let $\mathbb{C}(n)$ denote the quotient field of $\mathbb{C}[n]$, i.e., the field of all rational functions $P(n)/Q(n)$, where $P, Q \in \mathbb{C}[n]$. Suppose that $f \in \mathbb{C}^{\mathbb{N}}$ and $R \in \mathbb{C}(n)$. Then $f(n)R(n)$ is defined for n sufficiently large (i.e., when the denominator of $R(n)$ is nonzero). Thus we can define the germ $[f(n)R(n)] \in \mathcal{G}$ to be the germ of any function that agrees with $f(n)R(n)$ for n sufficiently large. It is easy to see that this definition of scalar multiplication makes \mathcal{G} into a vector space over the field $\mathbb{C}(n)$. We now come to the key characterization of P -recursive functions (or their germs).

12.17 Lemma. *A function $f \in \mathbb{C}^{\mathbb{N}}$ is P -recursive if and only if the vector space \mathcal{V}_f over $\mathbb{C}(n)$ spanned by the germs $[f(n)], [f(n+1)], [f(n+2)], \dots$ is finite-dimensional.*

Proof. Suppose that $f(n)$ satisfies equation (12.8). Let \mathcal{V}'_f be the vector space over $\mathbb{C}(n)$ spanned by $[f(n)], [f(n+1)], [f(n+2)], \dots, [f(n+d-1)]$, so $\dim_{\mathbb{C}(n)} \mathcal{V}'_f \leq d$. Equation (12.9) shows that $[f(n+d)] \in \mathcal{V}'_f$. Substitute $n+1$ for n in equation (12.9). We get that $[f(n+d+1)]$ is in the span (over $\mathbb{C}(n)$) of $[f(n+1)], [f(n+2)], \dots, [f(n+d)]$. Since these d germs are all in \mathcal{V}'_f , we get that $[f(n+d+1)] \in \mathcal{V}'_f$. Continuing in this way, we get by induction on k that $[f(n+d+k)] \in \mathcal{V}'_f$ for all $k \geq 0$, so $\mathcal{V}'_f = \mathcal{V}_f$. Thus \mathcal{V}_f is finite-dimensional.

Conversely, assume that $\dim_{\mathbb{C}(n)} \mathcal{V}_f < \infty$. Then for some d , the germs $[f(n)], [f(n+1)], \dots, [f(n+d)]$ are linearly dependent over $\mathbb{C}(n)$. Write down this linear dependence relation and clear denominators to get a recurrence (12.8) satisfied by f . Hence f is P -recursive. \square

We now have all the ingredients necessary for the main result of this section.

12.18 Theorem. *Let $f, g \in \mathcal{P}$ and $\alpha, \beta \in \mathbb{C}$. Then:*

- (a) $\alpha f + \beta g \in \mathcal{P}$
- (b) $fg \in \mathcal{P}$.

Proof. (a) By Lemma 12.17 it suffices to show that $\dim \mathcal{V}_{\alpha f + \beta g} < \infty$. Now by definition, the sum $\mathcal{V}_f + \mathcal{V}_g$ is the vector space consisting of all linear combinations $\gamma[u] + \delta[v]$, where $[u] \in \mathcal{V}_f$ and $[v] \in \mathcal{V}_g$ and $\gamma, \delta \in \mathbb{C}(n)$. In particular, $\mathcal{V}_f + \mathcal{V}_g$ contains all the germs $\alpha[f(n+k)] + \beta[g(n+k)] = [\alpha f(n+k)] + \beta g(n+k)$, $k \geq 0$. Hence

$$\mathcal{V}_{\alpha f + \beta g} \subseteq \mathcal{V}_f + \mathcal{V}_g.$$

Now if V and W are subspaces of some vector space, then $V + W$ is spanned by the union of a basis for V and basis for W . In particular, if V and W are finite-dimensional, then $\dim(V + W) \leq \dim V + \dim W$. Hence

$$\dim \mathcal{V}_{\alpha f + \beta g} \leq \dim(\mathcal{V}_f + \mathcal{V}_g) \leq \dim \mathcal{V}_f + \dim \mathcal{V}_g < \infty,$$

as was to be proved.

- (b) The proof is analogous to (a), except that instead of the sum $V + W$ we need the *tensor product* $V \otimes_K W$ over the field K . Recall from linear algebra that $V \otimes_K W$ may be thought of (somewhat naively) as the vector space spanned by all symbols $v \otimes w$, where $v \in V$ and $w \in W$, subject to the conditions

$$\begin{aligned}(v_1 + v_2) \otimes w &= v_1 \otimes w + v_2 \otimes w \\ v \otimes (w_1 + w_2) &= v \otimes w_1 + v \otimes w_2 \\ \alpha v \otimes w &= v \otimes \alpha w = \alpha(v \otimes w),\end{aligned}$$

where α is a scalar. A standard and simple consequence is that if V has the basis $\{v_1, \dots, v_m\}$ and W has the basis $\{w_1, \dots, w_n\}$, then $V \otimes_K W$ has the basis $v_i \otimes w_j$, for $1 \leq i \leq m$ and $1 \leq j \leq n$. In particular,

$$\dim(V \otimes_K W) = (\dim V)(\dim W).$$

Recall the basic “universality” property of the tensor product $V \otimes W = V \otimes_K W$: there is a bilinear map $\Psi: V \times W \rightarrow V \otimes W$ such that for any vector space Y and bilinear map $\Phi: V \times W \rightarrow Y$, there is a unique linear map $\varphi: V \otimes W \rightarrow Y$ for which $\Phi = \varphi\Psi$. In particular, there is a unique linear transformation $\varphi: \mathcal{V}_f \otimes_{\mathbb{C}(n)} \mathcal{V}_g \rightarrow \mathcal{G}$ satisfying

$$[f(n+i)] \otimes g[(n+j)] \xrightarrow{\varphi} [f(n+i)g(n+j)].$$

The image of φ contains all germs $[f(n+i)g(n+i)]$, so $\mathcal{V}_{fg} \subseteq \text{image}(\varphi)$. Thus

$$\dim \mathcal{V}_{fg} \leq \dim(\mathcal{V}_f \otimes_{\mathbb{C}(n)} \mathcal{V}_g) = (\dim \mathcal{V}_f)(\dim \mathcal{V}_g) < \infty,$$

and the proof follows. □

Notes for Chapter 12

The 100 prisoners problem was first considered by P. B. Miltersen. It appeared in a paper with A. Gál [44]. Further information on the history of this problem, together with a proof of optimality of the prisoners’ strategy, is given by E. Curtin and M. Warshauer [25].

The Oddtown theorem is due to E. R. Berlekamp [8]. Theorem 12.3 on decomposing K_n into complete bipartite subgraphs is due to R. L. Graham and H. O. Pollak [47][48]. For Fisher's original proof of the inequality $v \leq b$ for BIBD's and Bose's nonuniform generalization, see [35] and [12]. Sutner's original proof of the odd neighborhood theorem (Theorem 12.5) appears in [112], while the simpler proof of Y. Caro may be found in [20]. The odd neighborhood problem is also known as the *Lights Out Puzzle*. For a host of other applications of linear algebra along the lines of Sections 12.2–12.5, see the unpublished manuscript [5] of L. Babai and P. Frankl, and the book [78] of J. Matoušek.

The circulant Hadamard matrix conjecture was first mentioned in print by H. J. Ryser [97, p. 134], though its precise origin is obscure. The work of Turyn mentioned in the text appears in [116][117]. Some more recent progress is due to K. H. Leung and B. Schmidt [70].

While P -recursive functions and their cousins the D -finite series of Exercise 12.25 were known to 19th century analysts, the first systematic treatment of them did not appear until the paper of Stanley [104] in 1980, which includes a statement and proof of Theorem 12.18. For an exposition, see Stanley [110, §6.4].

Bibliography

- [1] T. van Aardenne-Ehrenfest and N. G. de Bruijn, Circuits and trees in oriented linear graphs, *Simon Stevin* (= *Bull. Belgian Math. Soc.*) **28** (1951), 203–217.
- [2] I. Anderson, *Combinatorics of Finite Sets*, Oxford University Press, Oxford-New York, 1987; corrected republication by Dover, New York, 2002.
- [3] G. E. Andrews, *The Theory of Partitions*, Addison-Wesley, Reading, Mass., 1976.
- [4] G. E. Andrews and K. Eriksson, *Integer Partitions*, Cambridge University Press, Cambridge, 2004.
- [5] L. Babai and P. Frankl, *Linear Algebra Methods in Combinatorics*, preliminary version 2 (September 1992), 216 pages.
- [6] E. A. Beem, Craige and Irene Schensted don't have a car in the world, *Maine Times*, March 12, 1982, pp. 20–21.
- [7] E. A. Bender and D. E. Knuth, Enumeration of plane partitions, *J. Combinatorial Theory* **13** (1972), 40–54.
- [8] E. R. Berlekamp, On subsets with intersections of even cardinality, *Canad. Math. Bull.* **12** (1969), 363–366.
- [9] H. Bidkhori and S. Kishore, Counting the spanning trees of a directed line graph, preprint, [arXiv:0910.3442](https://arxiv.org/abs/0910.3442).
- [10] E. D. Bolker, The finite Radon transform, in *Integral geometry* (*Brunswick, Maine, 1984*), *Contemp. Math.* **63**, Amer. Math. Soc., Providence, RI, 1987, pp. 27–50.

- [11] C. W. Borchardt, Ueber eine der Interpolation entsprechende Darstellung der Eliminations-Resultante, *J. reine angew. Math.* (= *Crelle's Journal*) **57** (1860), 111–121.
- [12] R. C. Bose, A note on Fisher's inequality for balanced incomplete block designs, *Ann. Math. Statist.* (1949), 619–620.
- [13] F. Brenti, Log-concave and unimodal sequences in algebra, combinatorics, and geometry: an update, in *Jerusalem Combinatorics '93*, *Contemp. Math.* **178**, Amer. Math. Soc., Providence, RI, 1994.
- [14] A. E. Brouwer and W. H. Haemers, *Spectra of Graphs*, Springer, New York, 2012.
- [15] N. G. de Bruijn, A combinatorial problem, *Proc. Koninklijke Nederlandse Akademie v. Wetenschappen* **49** (1946), 758–764, = *Indagationes Math.* **8** (1946), 461–467.
- [16] N. G. de Bruijn, Pólya's theory of counting, in *Applied Combinatorial Mathematics* (E. F. Beckenbach, ed.), Wiley, New York, 1964; reprinted by Krieger, Malabar, Florida, 1981.
- [17] N. G. de Bruijn, Acknowledgement of priority to C. Flye Sainte-Marie on the counting of circular arrangements of 2^n zeros and ones that show each n -letter word exactly once, Technische Hogeschool Eindhoven, T.H.-Report 75-WSK-06, 1975.
- [18] W. Burnside, *Theory of Groups of Finite Order*, Cambridge University Press, 1897.
- [19] W. Burnside, *Theory of Groups of Finite Order*, second edition, Cambridge University Press, 1911; reprinted by Dover, New York, 1955.
- [20] Y. Caro, Simple proofs to three parity theorems, *Ars Combin.* **42** (1996), 175–180.
- [21] N. Caspard, B. Leclerc, and B. Monjardet, *Finite Ordered Sets*, *Encyclopedia of Mathematics and Its Applications*, vol. 144, Cambridge University Press, Cambridge, 2012.

- [22] A. L. Cauchy, Mémoire sur diverses propriétés remarquables des substitutions régulières ou irrégulières, et des systèmes de substitutions conjuguées (suite), *C. R. Acad. Sci. Paris* **21** (1845), 972–987; *Oeuvres*, ser. 1, vol. 9, pp. 371–387.
- [23] A. Cayley, Note sur une formule pour la réversion des séries, *J. reine angew. Math.* (= *Crelle's Journal*) **52** (1856), 276–284.
- [24] A. Cayley, A theorem on trees, *Quart. J. Math.* **23** (1889), 376–378; *Collected Papers*, vol. 13, Cambridge University Press, 1897, pp. 26–28.
- [25] E. Curtin and M. Warshauer, The locker puzzle, *Math. Intelligencer* **28** (2006), 28–31.
- [26] D. M. Cvetković, M. Doob, and H. Sachs, *Spectra of Graphs: Theory and Applications*, 3rd ed., Johann Ambrosius Barth, Heidelberg/Leipzig, 1995.
- [27] D. M. Cvetković, P. Rowlinson, and S. Simić, *An Introduction to the Theory of Graph Spectra*, London Mathematical Society, Student Texts **75**, Cambridge University Press, Cambridge, 2010.
- [28] M. R. DeDeo and E. Velasquez, The Radon transform on \mathbb{Z}_n^k , *SIAM J. Discrete Math.* **18** (2004/05), 472–478 (electronic).
- [29] P. Diaconis and R. L. Graham, The Radon transform on \mathbb{Z}_2^k , *Pacific J. Math.* **118** (1985), 323–345.
- [30] P. Diaconis and R. L. Graham, *Magical Mathematics*, Princeton University Press, Princeton, 2012.
- [31] E. B. Dynkin, Some properties of the weight system of a linear representation of a semisimple Lie group (in Russian), *Dokl. Akad. Nauk SSSR (N.S)* **71** (1950), 221–224.
- [32] E. B. Dynkin, The maximal subgroups of the classical groups, in *Amer. Math. Soc. Transl., Ser. 2* **6** (1957), 245–378. Translated from *Trudy Moskov. Mat. Obsc.* **1**, 39–166.
- [33] K. Engel, *Sperner Theory*, Encyclopedia of Mathematics and Its Applications, no. 65, Cambridge University Press, Cambridge, 1997.

- [34] P. Fishburn, *Interval Orders and Interval Graphs: A Study of Partially Ordered Sets*, John Wiley & Sons, New York, 1985.
- [35] R. A. Fisher, An examination of the different possible solutions of a problem in incomplete blocks, *Ann. Eugenics* **10** (1940), 52–75.
- [36] C. Flye Sainte-Marie, Solution to question nr. 48, *l'Intermédiaire des Mathématiciens* **1** (1894), 107–110.
- [37] S. Fomin, Duality of graded graphs, *J. Algebraic Combinatorics* **3** (1994), 357–404.
- [38] S. Fomin, Schensted algorithms for dual graded graphs, *J. Algebraic Combinatorics* **4** (1995), 5–45.
- [39] J. S. Frame, G. de B. Robinson, and R. M. Thrall, The hook graphs of S_n , *Canad. J. Math.* **6** (1954), 316–324.
- [40] D. S. Franzblau and D. Zeilberger, A bijective proof of the hook-length formula, *J. Algorithms* **3** (1982), 317–343.
- [41] F. G. Frobenius, Über die Congruenz nach einem aus zwei endlichen Gruppen gebildeten Doppelmodul, *J. reine angew. Math. (= Crelle's Journal)* **101** (1887), 273–299; reprinted in *Gesammelte Abhandlungen*, vol. 2, Springer-Verlag, Heidelberg, 1988, pp. 304–330.
- [42] F. G. Frobenius, Über die Charaktere der symmetrischen Gruppe, *Sitzungsber. Kön. Preuss. Akad. Wissen. Berlin* (1900), 516–534; *Gesammelte Abh. III* (J.-P. Serre, ed.), Springer-Verlag, Berlin (1968), pp. 148–166.
- [43] W. E. Fulton, *Young Tableaux*, Student Texts **35**, London Mathematical Society, Cambridge University Press, Cambridge, 1997.
- [44] A. Gál and P. B. Miltersen, The cell probe complexity of succinct data structures, *Proc. 30th International Colloquium on Automata, Languages and Programming (ICALP)*, 2003, 332–344.
- [45] M. Gardner, Squaring the square, in *The 2nd Scientific American Book of Mathematical Puzzles and Diversions*, Simon and Schuster, New York, 1961.

- [46] I. J. Good, Normally recurring decimals, *J. London Math. Soc.* **21** (1947), 167–169.
- [47] R. L. Graham and H. O. Pollak, On the addressing problem for loop switching, *Bell System Tech. J.* **50** (1971), 2495–2519.
- [48] R. L. Graham and H. O. Pollak, On embedding graphs in squashed cubes, in *Lecture Notes in Mathematics* **303**, Springer Verlag, New York-Berlin-Heidelberg, 1973, pp. 99–110.
- [49] C. Greene, A. Nijenhuis, and H. S. Wilf, A probabilistic proof of a formula for the number of Young tableaux of a given shape, *Advances in Math.* **31** (1979), 104–109.
- [50] J. I. Hall, E. M. Palmer, and R. W. Robinson, Redfield’s lost paper in a modern context, *J. Graph Theory* **8** (1984), 225–240.
- [51] F. Harary and E. M. Palmer, *Graphical Enumeration*, Academic Press, New York/London, 1973.
- [52] F. Harary and R. W. Robinson, The rediscovery of Redfield’s papers, *J. Graph Theory* **8** (1984), 191–192.
- [53] G. H. Hardy, J. E. Littlewood, and G. Pólya, *Inequalities*, second ed., Cambridge University Press, Cambridge, England, 1952.
- [54] L. H. Harper, Morphisms for the strong Sperner property of Stanley and Griggs, *Linear and Multilinear Algebra* **16** (1984), 323–337.
- [55] T. W. Hawkins, The origins of the theory of group characters, *Arch. History Exact Sci.* **7** (1970/71), 142–170.
- [56] T. W. Hawkins, Hypercomplex numbers, Lie groups, and the creation of group representation theory, *Arch. History Exact Sci.* **8** (1971/72), 243–287.
- [57] T. W. Hawkins, New light on Frobenius’ creation of the theory of group characters, *Arch. History Exact Sci.* **12** (1974), 217–243.
- [58] A. P. Hillman and R. M. Grassl, Reverse plane partitions and tableaux hook numbers, *J. Combinatorial Theory (A)* **21** (1976), 216–221.

- [59] R. A. Horn and C. R. Johnson, *Matrix Analysis*, Cambridge University Press, Cambridge, 1985.
- [60] J. W. B. Hughes, Lie algebraic proofs of some theorems on partitions, in *Number Theory and Algebra* (H. Zassenhaus, ed.), Academic Press, New York, 1977, pp. 135–155.
- [61] A. Hurwitz, Über die Anzahl der Riemannschen Flächen mit gegebenen Verzweigungspunkten, *Math. Ann.* **55** (1902), 53–66.
- [62] A. Joyal, Une théorie combinatoire des séries formelles, *Advances in Math.* **42** (1981), 1–82.
- [63] N. D. Kazarinoff and R. Weitzenkamp, Squaring rectangles and squares, *Amer. Math. Monthly* **80** (1973), 877–888.
- [64] G. Kirchhoff, Über die Auflösung der Gleichungen, auf welche man bei der Untersuchung der Linearen Vertheilung galvanischer Ströme geführt wird, *Ann. Phys. Chem.* **72** (1847), 497–508.
- [65] C. Krattenthaler, Bijective proofs of the hook formulas for the number of standard Young tableaux, ordinary and shifted, *Electronic J. Combinatorics* **2** (1995), R13, 9pp.
- [66] D. E. Knuth, Permutations, matrices, and generalized Young tableaux, *Pacific J. Math.* **34** (1970), 709–727.
- [67] I. Krasikov and Y. Roditty, Balance equations for reconstruction problems, *Arch. Math. (Basel)* **48** (1987), 458–464.
- [68] J. P. S. Kung, Radon transforms in combinatorics and lattice theory, in *Combinatorics and ordered sets (Arcata, Calif., 1985)*, Contemp. Math. **57**, Amer. Math. Soc., Providence, RI, 1986, pp. 33–74.
- [69] M. A. A. van Leeuwen, The Robinson-Schensted and Schützenberger algorithms, Part 1: new combinatorial proofs, Preprint no. AM-R9208 1992, Centrum voor Wiskunde en Informatica, 1992.
- [70] K. H. Leung and B. Schmidt, New restrictions on possible orders of circulant Hadamard matrices, *Designs, Codes and Cryptography* **64** (2012), 143–151.

- [71] E. K. Lloyd, J. Howard Redfield: 1879–1944, *J. Graph Theory* **8** (1984), 195–203.
- [72] L. Lovász, A note on the line reconstruction problem, *J. Combin. Theory Ser. B* **13** (1972), 309–310.
- [73] L. Lovász, Random walks on graphs: a survey, in *Combinatorics. Paul Erdős is Eighty*, vol. 2, Bolyai Society Mathematical Studies **2**, Keszthely, Hungary, 1993, pp. 1–46.
- [74] D. Lubell, A short proof of Sperner’s lemma, *J. Combinatorial Theory* **1** (1966), 299.
- [75] P. A. MacMahon, Memoir on the theory of the partitions of numbers — Part I, *Phil. Trans. Royal Soc. London (A)* **187** (1897), 619–673; *Collected Works*, vol. 1 (G. E. Andrews, ed.), M.I.T. Press, Cambridge, Massachusetts, 1978, pp. 1026–1080.
- [76] P. A. MacMahon, Memoir on the theory of the partitions of numbers — Part IV, *Phil. Trans. Royal Soc. London (A)* **209** (1909), 153–175; *Collected Works*, vol. 1 (G. E. Andrews, ed.), M.I.T. Press, Cambridge, Massachusetts, 1978, pp. 1292–1314.
- [77] P. A. MacMahon, *Combinatory Analysis*, vols. 1 and 2, Cambridge University Press, Cambridge, 1915, 1916; reprinted in one volume by Chelsea, New York, 1960.
- [78] J. Matoušek, *Thirty-three Miniatures*, American Mathematical Society, Providence, RI, 2010.
- [79] J. W. Moon, *Counting Labelled Trees*, Canadian Math. Monographs, No. 1, Canadian Mathematical Congress, 1970.
- [80] V. Müller, The edge reconstruction hypothesis is true for graphs with more than $n \log_2 n$ edges, *J. Combin. Theory Ser. B* **22** (1977), 281–283.
- [81] P. M. Neumann, A lemma that is not Burnside’s, *Math. Sci.* **4** (1979), 133–141.

- [82] J.-C. Novelli, I. Pak, and A. V. Stoyanovskii, A new proof of the hook-length formula, *Discrete Math. and Theoretical Computer Sci.* **1** (1997), 053–067.
- [83] K. M. O’Hara, Unimodality of Gaussian coefficients: a constructive proof, *J. Combin. Theory Ser. A* **53** (1990), 29–52.
- [84] W. V. Parker, The matrices AB and BA , *Amer. Math. Monthly* **60** (1953), 316; reprinted in *Selected Papers on Algebra* (S. Montgomery et al., eds.), Mathematical Association of America, Washington, DC, pp. 331–332.
- [85] J. Pitman, Coalescent random forests, *J. Combinatorial Theory Ser. A* **85** (1999), 165–193,
- [86] G. Pólya, Kombinatorische Anzahlbestimmungen für Gruppen, Graphen und chemische Verbindungen, *Acta Math.* **68** (1937), 145–254.
- [87] G. Pólya and R. C. Read, *Combinatorial Enumeration of Groups, Graphs, and Chemical Compounds*, Springer-Verlag, New York/Berlin, 1987.
- [88] M. Pouzet, Application d’une propriété combinatoire des parties d’un ensemble aux groupes et aux relations, *Math. Zeit.* **150** (1976), 117–134.
- [89] M. Pouzet and I. G. Rosenberg, Sperner properties for groups and relations, *European J. Combin.* **7** (1986), 349–370.
- [90] R. A. Proctor, A solution of two difficult combinatorial problems with linear algebra, *Amer. Math. Monthly* **89** (1982), 721–734.
- [91] H. Prüfer, Neuer Beweis eines Satzes über Permutationen, *Arch. Math. Phys.* **27** (1918), 742–744.
- [92] J. Radon, Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten, *Berichte über die Verhandlungen der Königlich Sächsischen Gesellschaft der Wissenschaften zu Leipzig* **69** (1917), 262–277; translation by P. C. Parks, On the determination of functions from their integral values along certain manifolds, *IEEE Transactions on Medical Imaging* **5** (1986), 170–176.

- [93] J. H. Redfield, The theory of group reduced distributions, *American J. Math.* **49** (1927), 433–455.
- [94] J. H. Redfield, Enumeration by frame group and range groups, *J. Graph Theory* **8** (1984), 205–223.
- [95] J. B. Remmel, Bijective proofs of formulae for the number of standard Young tableaux, *Linear and Multilinear Algebra* **11** (1982), 45–100.
- [96] G. de B. Robinson, On the representations of S_n , *Amer. J. Math.* **60** (1938), 745–760.
- [97] H. J. Ryser, *Combinatorial Mathematics*, Mathematical Association of America, Washington, DC, 1963.
- [98] B. E. Sagan, *The Symmetric Group*, second ed., Springer, New York, 2001.
- [99] C. E. Schensted, Longest increasing and decreasing subsequences, *Canadian J. Math.* **13** (1961), 179–191.
- [100] J. Schmid, A remark on characteristic polynomials, *Amer. Math. Monthly* **77** (1970), 998–999; reprinted in *Selected Papers on Algebra* (S. Montgomery et al., eds.), Mathematical Association of America, Washington, DC, pp. 332–333.
- [101] C. A. B. Smith and W. T. Tutte, On unicursal paths in a network of degree 4, *Amer. Math. Monthly* **48** (1941), 233–237.
- [102] E. Sperner, Ein Satz über Untermengen einer endlichen Menge, *Math. Z.* **27**(1) (1928), 544–548.
- [103] R. Stanley, Weyl groups, the hard Lefschetz theorem, and the Sperner property, *SIAM J. Algebraic and Discrete Methods* **1** (1980), 168–184.
- [104] R. Stanley, Differentiably finite power series, *European J. Combinatorics* **1** (1980), 175–188.
- [105] R. Stanley, Quotients of Peck posets, *Order* **1** (1984), 29–34.
- [106] R. Stanley, Differential posets, *J. Amer. Math. Soc.* **1** (1988), 919–961.

- [107] R. Stanley, Unimodal and log-concave sequences in algebra, combinatorics, and geometry, in *Graph Theory and Its Applications: East and West*, Ann. New York Acad. Sci., vol. 576, 1989, pp. 500–535.
- [108] R. Stanley, Variations on differential posets, in *Invariant Theory and Tableaux* (D. Stanton, ed.), The IMA Volumes in Mathematics and Its Applications, vol. 19, Springer-Verlag, New York, 1990, pp. 145–165.
- [109] R. Stanley, *Enumerative Combinatorics*, vol. 1, 2nd ed., Cambridge University Press, Cambridge, 2012.
- [110] R. Stanley, *Enumerative Combinatorics*, vol. 2, Cambridge University Press, New York/Cambridge, 1999.
- [111] B. Stevens, G. Hurlbert, and B. Jackson, eds., Special issue on “Generalizations of de Bruijn cycles and Gray codes,” *Discrete Math.* **309** (2009).
- [112] K. Sutner, Linear cellular automata and the Garden-of-Eden, *Math. Intelligencer* **11** (1989), 49–53.
- [113] J. J. Sylvester, On the change of systems of independent variables, *Quart. J. Math.* **1** (1857), 42–56; *Collected Mathematical Papers*, vol. 2, Cambridge, 1908, pp. 65–85.
- [114] J. J. Sylvester, Proof of the hitherto undemonstrated fundamental theorem of invariants, *Phil. Mag.* **5** (1878), 178–188; *Collected Mathematical Papers*, vol. 3, Chelsea, New York, 1973, pp. 117–126.
- [115] W. T. Trotter, *Combinatorics and Partially Ordered Sets: Dimension Theory*, Johns Hopkins Studies in the Mathematical Sciences **6**, Johns Hopkins University Press, Baltimore, 1992.
- [116] R. J. Turyn, Character sums and difference sets, *Pacific J. Math.* **15** (1965), 319–346.
- [117] R. J. Turyn, Sequences with small correlation, in *Error Correcting Codes* (H. B. Mann, ed.), Wiley, New York, 1969, pp. 195–228.
- [118] W. T. Tutte, The dissection of equilateral triangles into equilateral triangles, *Proc. Cambridge Phil. Soc.* **44** (1948), 463–482.

- [119] W. T. Tutte, Lectures on matroids, *J. Res. Nat. Bur. Standards Sect. B* **69** (1965), 1–47.
- [120] W. T. Tutte, The quest of the perfect square, *Amer. Math. Monthly* **72** (1965), 29–35.
- [121] E. M. Wright, Burnside’s lemma: a historical note, *J. Combinatorial Theory (B)* **30** (1981), 89–90.
- [122] A. Young, Qualitative substitutional analysis (third paper), *Proc. London Math. Soc. (2)* **28** (1927), 255–292.
- [123] D. Zeilberger, Kathy O’Hara’s constructive proof of the unimodality of the Gaussian polynomials, *Amer. Math. Monthly* **96** (1989), 590–602

Index

- 4483130665195087, 46, 57
- van Aardenne-Ehrenfest, Tanya, 198, 199
- access time, 35
- acts on (by a group), 59
- acyclic (set of edges), 210
- adjacency matrix, 10
- adjacent (vertices), 9
- adjoint (operator), 51
- algebraic (Laurent series), 255
- algebraic integer, 243
- Anderson, Ian, 56
- Andrews, George W. Eyre, 95
- antichain, 48
- antisymmetry, 45
- automorphism
 - of a graph, 126
 - of a poset, 60
- automorphism group
 - of a graph, 126
 - of a poset, 61
- Babai, László, 250
- balanced digraph, 190
- balanced incomplete block design, 236
- basis matrix, 209
- Bender, Edward Anton, 158
- Berlekamp, Elwyn Ralph, 250
- Bernardi, Olivier, 179
- BEST theorem, 198
- BIBD, 236
- Bidkhor, Hoda, 199
- binary de Bruijn sequence, *see* de Bruijn sequence
- binary sequence, 195
- Binet-Cauchy theorem, 170, 171
- binomial moment, 43
- bipartite graph, 17
- bipartition, 17
- block (of a block design), 236
- block design, 236
- Bolker, Ethan David, 28
- bond, 205
- bond space, 205
- boolean algebra, 45
- Borchardt, Carl Wilhelm, 184
- Bose, Raj Chandra, 236, 250
- Brenti, Francesco, 71
- de Bruijn, Nicolaas Govert, 125, 198, 199
- de Bruijn sequence, 195
- de Bruijn graph, 196
- bump
 - in RSK algorithm for CSPP, 149
 - in RSK algorithm for SYT, 144
- Burnside's lemma, 105, 124
- Burnside, William, 124
- \mathbb{C} -algebra, 246
- Caro, Yair, 250
- Caspard, Nathalie, 55

- Cauchy, Augustin Louis, 124
- Cauchy-Binet theorem, 170, 171
- Cauchy-Frobenius lemma, 105, 124
- Cayley, Arthur, 184
- chain (in a poset), 46
- characteristic polynomial, 17
- characteristic vector (of a set), 234
- circuit, 204
- circulant (matrix), 240
- circulation, 203
- closed walk, 12
- coboundary, 205
- Collatz, Lothar, 16
- coloring, 100
- column-strict plane partition, 148
- commutative diagram, 65
- complementary graph, 186
- complete bipartite graph, 17, 185
- complete graph, 13
- complete p -partite graph, 18
- complexity (of a graph), 170
- conjugate
 - of an algebraic number, 244
 - partition, 80
- connected graph, 169
- cotree, 212
- covers (in a poset), 46
- CSPP, 148
- cube (graph), 21
- Curtin, Eugene, 249
- Cvetković, Dragoš M., 16
- cycle index polynomial, 108
- cycle indicator, 124
 - of a group of permutations, 108
 - of a permutation, 108
- cycle space, 204
- cyclic group, 30
- cyclomatic number, 211
- cyclotomic polynomial, 245
- Dedekind, Julius Wilhelm Richard, 28
- DeDeo, Michelle Rose, 28
- degree (of a vertex), 18, 33
- deleted neighborhood (of a vertex), 252
- D -finite, 255
- Diaconis, Persi Warren, 28, 199
- diagonal (of a power series), 255
- diagram (of a partition), 78
- differentiably finite, 255
- digraph, 189
 - of a permutation, 181
- dihedral necklace, 112
- direct product (of posets), 97
- directed graph, 189
- Doob, Michael, 16
- doubly-rooted tree, 180
- down (linear transformation), 51
- dual (of a planar graph), 221
- Dynkin, Eugene (Evgenii) Borisovitch, 95
- edge reconstruction conjecture, 68
- edge set (of a graph), 9
- eigenvalues of a graph, 11
- elementary cycle, 204
- Engel, Konrad, 56
- equivalent colorings, 100, 101
- Erdős-Moser conjecture, 94
 - weak, 94
- Eriksson, Kimmo, 95
- Euler phi-function, 111
- Euler's constant, 233
- Eulerian cycle (in a graph), 189
- Eulerian digraph, 190
- Eulerian graph, 189

- Eulerian tour
 - in a digraph, 189
 - in a graph, 189
- extended Smith diagram, 226
- face (of a planar embedding), 220
- faithful action, 59
- Fermat's Last Theorem, 242
- Ferrers diagram, 78
- Fibonacci number, 165, 185, 246
- final vertex (of an edge)
 - in a digraph, 189
 - in an orientation, 173
- Fishburn, Peter, 55
- Fisher, Ronald Aylmer, 236, 250
- Flye Sainte-Marie, Camille, 199
- Fomin, Sergey Vladimirovich, 157
- forest, 182, 211
- Forsyth, Andrew Russell, 158
- Frame, James Sutherland, 156
- Frame–Robinson–Thrall, 133
- Frankl, Peter, 250
- Franzblau, Deborah Sharon, 157
- Frobenius, Ferdinand Georg, 28, 124, 125, 156
- Fulton, William Edgar, 158
- Fundamental Theorem of Algebra, 70
- Gál, Anna, 249
- Gardner, Martin, 228
- Gauss' lemma, 241
- Gaussian coefficient, 82
- generalized ballot sequence, 156
- generating function, 15, 104
- G -equivalent, 60
 - colorings, 101
- germ, 247
- Good, Irving John, 198
- graded poset, 47
- Graham, Ronald Lewis, 28, 199, 250
- graph, 9
- Grassl, Richard, 156
- Greene, Curtis, 156
- group determinant, 28
- group reduction function, 125
- Hadamard matrix, 240
- Hamiltonian cycle, 30
- Hamming weight, 25
- Harary, Frank, 125
- Hardy, Godfrey Harold, 71
- Harper, Lawrence Hueston, 71
- Hasse diagram, 46
- Hasse walk, 132
- Hawkins, Thomas W., 29
- Hillman, Abraham, 156
- hitting time, 35
- hook length formula, 133, 156
- Horn, Roger Alan, 40
- Hughes, J. W. B., 95
- Hurwitz, Adolf, 125
- incidence matrix
 - Oddtown, 234
 - of a digraph, 207
 - of a graph, 173
- incident, 9
- indegree (of a vertex), 190
- induced subgraph, 239
- initial vertex (of an edge)
 - in a digraph, 189
 - in an orientation, 173
- internal zero, 69
- inverse bump (in RSK algorithm), 146
- isolated vertex, 36
- isomorphic

- graphs, 66
- posets, 46
- isomorphism class (of simple graphs), 66
- Johnson, Charles Royal, 40
- Joyal, André, 180, 184
- Kazarinoff, Nicholas D., 228
- Kirchhoff's laws, 214
- Kirchhoff, Gustav Robert, 184, 228
- Kishore, Shaunak, 199
- Knuth, Donald Ervin, 139, 157, 158
- Krasikov, Ilia, 71
- Krattenthaler, Christian Friedrich, 157
- Kronecker, Leopold, 244
- Kummer, Ernst Eduard, 242
- Kung, Joseph PeeSin, 28
- laplacian matrix, 174
- lattice, 77
- lattice permutation, 156
- Laurent series, 254
- Leclerc, Bruno, 55
- van Leeuwen, Marc A. A., 157
- length
 - of a chain, 47
 - of a necklace, 72, 111
 - of a walk, 10
- Leung, Ka Hin, 250
- level (of a ranked poset), 47
- Lights Out Puzzle, 250
- Littlewood, Dudley Ernest, 157
- Littlewood, John Edensor, 71
- Littlewood–Richardson rule, 157
- log-concave, 69
 - polynomial, 75
- logarithmically concave, 69
- loop
 - in a digraph, 190
 - in a graph, 9
- Lovász, László, 40, 71
- Lubell, David, 48, 53, 56
- MacMahon, Percy Alexander, 156–158
- mail carrier, 195
- Markov chain, 33
- Matoušek, Jiří, 250
- matrix
 - irreducible, 35
 - nonnegative, 35
 - permutation, 36
- matrix analysis, 40
- Matrix-Tree Theorem, 176, 184
- matroid, 223
- maximal chain, 47
- Miltersen, Peter Bro, 249
- Möbius function, 127
- Monjardet, Bernard, 55
- Moon, John W., 184
- Müller, Vladimír, 71
- multiple edge, 10
- multiset, 9
- n -cycle, 31
- necklace, 72, 111
- neighborhood (of a vertex), 238
- Newton, Isaac, 69, 71
- Nijenhuis, Albert, 156
- N-matrix, 149
- no internal zero, 69
- nonuniform Fisher inequality, 237
- Novelli, Jean-Christophe, 157
- Oddtown, 233
- O'Hara, Kathleen Marie, 87, 95
- Ohm's law, 215
- orbit, 60

- order (of a poset), 46
- order-matching, 49
 - explicit for B_n , 54
- order-raising operator, 50
- orientation (of a graph), 173
- orthogonal complement, 208
- orthogonal Lie algebra, 95
- orthogonal subspace, 239
- outdegree (of a vertex), 190

- Pak, Igor M., 157
- Palmer, Edgar Milan, 125
- parallel connection, 215
- Parker, William Vann, 95
- part (of a partition of n), 77
- partially ordered set, 45
- partition
 - of a set X , 60
 - of an integer n , 77
- Pascal's triangle, 47, 83
 - q -analogue, 83
- path (in a graph), 170
 - closed, 30
- perfect squared rectangle, 223
- Perron–Frobenius theorem, 36
- physical intuition, 215
- Pitman, James William, 182, 184
- planar embedding, 220
- planar graph, 220
- plane partition, 147
 - history of, 157
- planted forest, 182
- pole (of a Smith diagram), 224
- Pollak, Henry Otto, 250
- Pólya, George (György), 71, 99, 125
- Pólya theory, 99
- polynomially recursive function, 246
- poset, 45
 - positive definite, 237
 - positive semidefinite, 52, 237
 - potential, 205
 - potential difference, 205
- Pouzet, Maurice André, 56, 71
- P -recursive function, 246
- primitive necklace, 126
- probability matrix, 33, 41
- Proctor, Robert Alan, 95
- Prüfer sequence, 180
- Prüfer, Ernst Paul Heinz, 180, 184

- q -binomial coefficient, 57, 82
- quantum order-matching, 51
- quotient poset, 61

- Radon transform, 23
- Radon, Johann Karl August, 28
- rank
 - of a boolean algebra, 45
 - of a graded poset, 47
 - of a poset element, 47
- rank-generating function, 47
- rank-symmetric, 47
- rank-unimodal, 47
- reciprocity theorem, 128
- Redfield, John Howard, 99, 125
- reduced incidence matrix, 175
- reflexivity, 45
- region (of a planar embedding), 220
- regular graph, 34
- Remmel, Jeffrey Brian, 157
- Robinson, Gilbert de Beauregard, 139, 156, 157
- Robinson–Schensted correspondence, 139
- Roditty, Yehuda, 71
- Rolle's theorem, 70

- root (of a tree), 182
- Rosenberg, Ivo G., 56, 71
- row insertion, 144, 149
- Rowlinson, Peter, 16
- RSK algorithm, 139, 157
- Ryser, Herbert John, 250

- Sachs, Horst, 16
- Sagan, Bruce, Eli, 158
- Schensted, Craige Eugene, 139, 157
- Schmid, Josef, 95
- Schmidt, Bernard, 250
- Schützenberger, Marcel-Paul, 157
- semidefinite, 52, 237
- series connection, 215
- series-parallel network, 216
- shape (of a CSPP), 148
- shifted Ferrers diagram, 152
- Simić, Slobodan, 16
- simple (squared square), 227
- simple graph, 10
- simple group, 128
- Sinogowitz, Ulrich, 16
- Smith diagram, 224
- Smith, Cedric Austen Bardell, 198
- solid partition, 155
- spanning subgraph, 66, 170
- spectral graph theory, 16
- Sperner poset, 48
- Sperner property, 48
- Sperner's theorem, 48
- Sperner, Emanuel, 56
- squared rectangle, 223
- stabilizer, 63
- standard Young tableau, 133
- Stanley, Richard Peter, 56, 71, 95, 125, 157, 158, 199, 250
- stationary distribution, 41

- von Staudt, Karl Georg Christian, 184
- Stirling number, signless of the first kind, 114
- Stoyanovskii, Alexander V., 157
- strongly log-concave, 69
- sum (of vector spaces), 248
- support (of a function), 209
- Sutner, Klaus, 238, 250
- switching (at a vertex), 74
- switching reconstructible, 74
- Sylvester, James Joseph, 87, 95, 184
- symmetric chain decomposition, 55
- symmetric function, 104
- symmetric plane partition, 167
- symmetric sequence, 47
- SYT, 133

- tensor product (of vector spaces), 249
- Thrall, Robert McDowell, 156
- total resistance, 215
- totient function, 111
- tour (in a digraph), 189
- trace (of a plane partition), 167
- transitive (group action), 74, 105
- transitivity, 45
- transport, 65
- transposition, 128
- tree, 169
- Trotter, William Thomas, Jr., 56
- Turyn, Richard Joseph, 241, 250
- Tutte, William Thomas, 198, 199, 228
- two-line array, 149
- type
 - of a Hasse walk, 132
 - of a permutation, 108
- unimodal sequence, 47
- unimodular (matrix), 213

- universal cycle for \mathfrak{S}_n , 202
- universality (of tensor products), 249
- up (linear transformation), 50

- valid λ -word, 135
- Velasquez, Elinor Laura, 28
- Venkataramana, Praveen, 185
- vertex bipartition, 17
- vertex reconstruction conjecture, 67
- vertex set, 9

- walk, 10
- Warshauer, Max, 249
- weakly switching-reconstructible, 74
- weight
 - of a binary vector, 25
 - of a necklace, 72
- Weitzenkamp, Roger, 228
- Wheatstone bridge, 216
- Wilf, Herbert Saul, 156
- wreath product, 85

- Young diagram, 78
- Young, Alfred, 77, 156
- Young's lattice, 77

- Zeilberger, Doron, 95, 157
- Zyklus, 108