Chapter 10

Enumeration under group action

We might well have titled this chapter *Pólya–de Bruijn theory*, since its foundations are due to George Pólya (1937) and Nicolaas de Bruijn (1959). This theory furnishes several powerful tools for enumerating equivalence classes of discrete structures representable by functions from one finite set to another. In addition to the now familiar examples of words and distributions, any relation $R \subseteq A \times B$ can be so represented by its characteristic function $\chi_R : A \times B \to \{0,1\}$. So these tools are applicable to a fairly broad class of discrete structures. Since the domain and codomain of a function are sets, however, their members are by definition distinct (equivalently, *distinguishable*, or *labeled*). But we have already encountered enumeration problems in which the relevant entities (for example, balls or boxes) are postulated to be unlabeled. Readers who have completed Exercise 2.9 on *domain equivalence* and Exercise 6.2 on *codomain equivalence* will have seen a preview of how to model total indistinguishability of members of the domain or codomain of a function. The techniques described below allow one to define (and enumerate) equivalence classes of functions based on a notion of partial indistinguishability, defined by any permutation group on their domains or codomains. The algebraic prerequisites for this chapter are more substantial than those of previous chapters and include familiarity with permutation groups and the cycle decomposition of a permutation. See Warner (1965, Section 8) for a brief review of these topics.

10.1. Permutation groups and orbits

If X is a nonempty finite set, the set of all permutations g of X (i.e., bijections $g : X \to X$), denoted $S(X)$, is a group under the operation of composition of functions, known as the *symmetric group on X*. Any subgroup G of $S(X)$ is called a *permutation group on X*. Given such a permutation group and $x, y \in X$, we say that x and y are *congruent* mod$G$, symbolized $x \equiv y(G)$, if there exists a permutation $g \in G$ such that $g(x) = y$.

Congruence mod G is clearly an equivalence relation on X, and the equivalence classes induced by this equivalence relation are called *orbits of* G.

Example 1. The permutation group $G = S(X)$ has the single orbit X, since for all $x, y \in X$, there exists a permutation $g \in S(X)$ such that $g(x) = y$.

Example 2. If G consists solely of the identity permutation i_X, the orbits of G are (all of) the singleton subsets of X.

Example 3. If G is the permutation group on $X = [6]$ generated by $g = (1)(23)(465)$, the orbits of G are simply $\{1\}, \{2, 3\}$, and $\{4, 5, 6\}$. More generally, the orbits of any cyclic permutation group correspond to the cycles of the generator of that group. Indeed, the orbits of a permutation group are usefully construed as generalizations of the cycles of a permutation.

The following notation will be used in the remainder of this chapter. If g is a permutation of a finite set X and j is a positive integer, then $\lambda_j(g)$ denotes the number of cycles of g of length j. In particular, $\lambda_1(g) =$ the number of fixed points of g. If G is a permutation group on X, then $O(G)$ denotes the set of orbits of G. The entire theory of enumeration under group action is based on our first theorem, which furnishes a formula for the number of orbits of an arbitrary permutation group.

Theorem 10.1.1 (Burnside's lemma)**.** *If G is a permutation group on X, then the number of orbits of G is equal to the average number of fixed points, taken over all permutations $g \in G$, that is,*

$$|O(G)| = \frac{1}{|G|} \sum_{g \in G} \lambda_1(g). \tag{10.1.1}$$

Proof. For each $x \in X$, let O_x denote the orbit of G containing x, and let $G_x := \{g \in G : g(x) = x\}$. G_x is a subgroup of G, called the *stabilizer of x in G*. We first show that for all $x \in X$,

$$|G| = |G_x| \cdot |O_x|. \tag{10.1.2}$$

The simplest proof of this identity is based on the existence of a bijection from O_x to G/G_x, the set of (say) right cosets of G_x. In what follows, however, we establish (10.1.2) by an elementary argument that makes no use of coset decompostions of G. Our proof is based on the fact that the map $\psi : G \to O_x$ defined by $\psi(g) = g(x)$, is a $|G_x|$-to-one surjection, i.e., that

$$|\psi^{-1}(\{y\})| = |G_x|, \text{ for all } y \in O_x. \tag{10.1.3}$$

If $y = x$, (10.1.3) is obvious, since $\psi^{-1}(\{x\}) = \{g \in G : g(x) = x\} = G_x$. More generally, for any $y \in O_x$, choose $h \in G$ such that $h(y) = x$. Then $g \mapsto h \circ g$ is a bijection from $\psi^{-1}(\{y\}) = \{g \in G : g(x) = y\}$ to G_x. By (10.1.2) we then have

$$\sum_{g \in G} \lambda_1(g) = \sum_{g \in G} \sum_{\substack{x \in X \\ g(x) = x}} 1 = \sum_{x \in X} \sum_{\substack{g \in G \\ g(x) = x}} 1 = \sum_{x \in X} |G_x| = |G| \sum_{x \in X} \frac{1}{|O_x|}$$

$$= |G| \sum_{O \in O(G)} \sum_{\substack{x \in X \\ O_x = O}} \frac{1}{|O|} = |G| \sum_{O \in O(G)} \frac{|O|}{|O|} = |G| \cdot |O(G)|. \qquad \square$$

Remark 10.1.2 (Some elementary applications of Burnside's lemma)**.**

(i) We observed above that the symmetric group $S(X)$ has a single orbit. By (10.1.1) it follows that the average number of fixed points, taken over all permutations of X, is equal to 1, a result that can also be derived directly from formula (3.3.11).

(ii) If $G = i_X$, then, as observed above, the orbits of G are all of the singleton subsets of X, and so $|O(G)| = |X|$ equals the average number of fixed points, taken over all permutations in G which equals the number of fixed points of the sole element i_X of G.

(iii) Let X be the set of all $n!$ visually distinct circular arrangements of n distinct objects, with G comprising the set of all clockwise rotations of such arrangements by $2\pi j/n$, for $j = 0, \ldots, n-1$. Rotational equivalence classes of such arrangements coincide here with orbits of G, and so the number of rotational equivalence classes of n distinct objects is equal to $\frac{1}{n}(n! + 0 + \cdots + 0) = (n-1)!$, a well-known result with an elementary proof.

(iv) Let X be the set of all binary words of length 10, and let G be the two-element permutation group on X consisting of the identity permutation and the permutation $g(b_1, b_2, \ldots, b_{n-1}, b_n) = (b_n, b_{n-1}, \ldots, b_2, b_1)$. By (10.1.1) the number of orbits of G is equal to $\frac{1}{2}(2^{10} + 2^5) = 2^9 + 2^4$, since the fixed points of g are precisely the palindromes. Alternatively, the number of singleton orbits of G is equal to the number of palindromes, namely, 2^5, and the number of doubleton orbits is equal to half the number of nonpalindromes, namely, $((2^{10} - 2^5)/2 = 2^9 - 2^4$, and so the total number of orbits is equal to $2^9 + 2^5 - 2^4 = 2^9 + 2^4$.

One can also use Burnside's lemma to derive formula (8.2.2) for the number of rotational equivalence classes of words of length n in an alphabet of cardinality k. But in this case, as well as many others in which the goal is to count the orbits of a permutation group G on X, the members of X are actually *functions* from a domain D to a codomain C, and the permutation group on this set of functions is induced by a permutation group on D. In what follows, we take this observation into explicit account, with significant gains in clarity.

10.2. Pólya's first theorem

In what follows, $|D| = n$ and $|C| = k$, where, to avoid trivialities, $n, k \geq 2$. For vividness, the elements of C are sometimes thought of as a set of *colors*, so that each function $\varphi : D \to C$ corresponds to a *coloring* of the members of the domain D. Associated with each permutation group G on D is a binary relation $\sim$ on C^D, where $\varphi_1 \sim \varphi_2$ if and only if there exists a permutation $g \in G$ such that $\varphi_1 \circ g = \varphi_2$. It is easy to see that $\sim$ is an equivalence relation on C^D.

The equivalence classes induced by $\sim$ are called *G-equivalence classes* (abbreviated in what follows as *G-classes*). In order to apply Burnside's lemma to determine the number of G -classes, we need to identify such classes with the orbits of a permutation group (we will call it $\overline{G}$) on C^D. $\overline{G}$ is constructed from G as follows. For each $g \in G$ and each $\varphi \in C^D$, let $\overline{g}(\varphi) := \varphi \circ g$. Clearly, $\overline{g} : C^D \to C^D$. Moreover, $\overline{g}$ is injective (and

hence a permutation of the finite set C^D). For if $\overline{g}(\varphi_1) = \overline{g}(\varphi_2)$, then $\varphi_1 \circ g = \varphi_2 \circ g$ and so $\varphi_1 = \varphi_2 \circ g \circ g^{-1} = \varphi_2$. Let $\overline{G} := \{\overline{g} : g \in G\}$.

Lemma 10.2.1. *$\overline{G}$ is a permutation group on C^D, and $|\overline{G}| = |G|$. The orbits of $\overline{G}$ are identical with G-classes.*

Proof. To show that $\overline{G}$ is a subgroup of the symmetric group $S(C^D)$ it suffices to show that $\overline{g}_1 \circ \overline{g}_2 \in \overline{G}$ whenever $\overline{g_1}, \overline{g_2} \in \overline{G}$. But for every $\varphi \in C^D$,

$$\overline{g_1} \circ \overline{g_2}(\varphi) = (\varphi \circ g_2) \circ g_1 = \varphi \circ (g_2 \circ g_1) = \overline{g_2 \circ g_1}(\varphi),$$

and so $\overline{g_1} \circ \overline{g_2} = \overline{g_2 \circ g_1} \in \overline{G}$ (i.e., the map $g \mapsto \overline{g}$ is an *antihomomorphism*). Since $\overline{G}$ is defined as the range of the map $g \mapsto \overline{g}$, to show that $|\overline{G}| = |G|$, it suffices to show that this map is injective. So suppose that $g_1, g_2 \in G$ and $g_1 \neq g_2$. Then there exists $x \in D$ such that $g_1(x) \neq g_2(x)$. Since $k \geq 2$, there exists $\varphi \in C^D$ such that $\varphi(g_1(x)) \neq \varphi(g_2(x))$, whence $\varphi \circ g_1 \neq \varphi \circ g_2$, i.e., $\overline{g_1}(\varphi) \neq \overline{g_2}(\varphi)$. Hence, $\overline{g_1} \neq \overline{g_2}$. Finally, suppose that $\varphi_1, \varphi_2 \in C^D$. Then $\varphi_1 \sim \varphi_2 \Leftrightarrow$ there exists $g \in G$ such that $\varphi_1 \circ g = \varphi_2 \Leftrightarrow \overline{g}(\varphi_1) = \varphi_2 \Leftrightarrow \varphi_1 \equiv \varphi_2(\overline{G})$. □

The formula for the number of G-classes involves the important polynomial $Z(G : z_1, \ldots, z_n)$, known as the *cycle index* (in German, *Zyklenzeiger*) of G, and defined by

$$Z(G; z_1, \ldots, z_n) = \frac{1}{|G|} \sum_{g \in G} z_1^{\lambda_1(g)} \cdots z_n^{\lambda_n(g)}, \tag{10.2.1}$$

where, as noted earlier, $\lambda_j(g)$ is equal to the number of cycles of g of length j. The cycle index is actually a probability generating function, for grouping like terms in (10.2.1) yields the formula

$$Z(G; z_1, \ldots, z_n) = \sum_{\substack{\lambda_1 + \cdots + n\lambda_n = n \\ \lambda_j \geq 0}} (|\{g \in G \text{ with cycle structure } 1^{\lambda_1} \cdots n^{\lambda_n}\}|/|G|) z_1^{\lambda_1} \cdots z_n^{\lambda_n}. \tag{10.2.2}$$

In the above formula, the coefficient of $z_1^{\lambda_1} \cdots z_n^{\lambda_n}$ is equal to the probability that a permutation randomly chosen from the permutation group G has the cycle structure $1^{\lambda_1} \cdots n^{\lambda_n}$.

Theorem 10.2.2 (Pólya's first theorem)**.** *The number of G-classes into which C^D is partitioned by $\sim$ is equal to $Z(G; k, \ldots, k)$.*

Proof. By Lemma 10.2.1, along with Theorem 10.1.1, the number of G-classes is equal to

$$\begin{aligned} |O(\overline{G})| &= \frac{1}{|\overline{G}|} \sum_{g \in \overline{G}} \lambda_1(\overline{g}) \\ &= \frac{1}{|G|} \sum_{g \in G} |\{\varphi \in C^D : \overline{g}(\varphi) = \varphi \circ g = \varphi\}| = \frac{1}{|G|} \sum_{g \in G} k^{\lambda_1(g) + \cdots + \lambda_n(g)}, \end{aligned}$$

since $\varphi \circ g = \varphi$ if and only if φ is constant on the cycles of g, and g has $\lambda_1(g) + \cdots + \lambda_n(g)$ cycles. □

In view of the above theorem, it would be useful to have a list of the cycle indices of various permutation groups. A fairly extensive such list appears in the text *Graphical Enumeration* (Harary and Palmer 1973). Nevertheless, let us note a few basic results, taking $D = [n]$ in the remainder of this section.

(i) If $i_{[n]}$ denotes the identity permutation on $[n]$, then clearly,

$$Z(\{i_{[n]}\}; z_1, \ldots, z_n) = z_1^n. \tag{10.2.3}$$

It follows from Pólya's first theorem that the number of $\{i_{[n]}\}$-classes is equal to k^n, in accord with the obvious fact that the $\{i_{[n]}\}$-classes are all singletons.

(ii) Let S_n denote the symmetric group on $[n]$.

Theorem 10.2.3.

$$Z(S_n; z_1, \ldots, z_n) = \sum_{\lambda_1 + \cdots + n\lambda_n = n} \frac{1}{1^{\lambda_1} \cdots n^{\lambda_n} \lambda_1! \cdots \lambda_n!} z_1^{\lambda_1} \cdots z_n^{\lambda_n}. \tag{10.2.4}$$

Proof. Immediate, from Cauchy's formula (Theorem 6.6.6). □

(iii) Let C_n be the cyclic permutation group on $[n]$ generated by the cycle $g = (12 \cdots n)$.

Theorem 10.2.4.

$$Z(C_n; z_1, \ldots, z_n) = \frac{1}{n} \sum_{d|n} \phi(n/d) z_{n/d}^d = \frac{1}{n} \sum_{d|n} \phi(d) z_d^{n/d}. \tag{10.2.5}$$

Proof. Recall that if $m \in \mathbb{Z}$ and $n \in \mathbb{P}$, $m \bmod n$ denotes the unique $r \in \{0, 1, \ldots, n-1\}$ such that $m \equiv r \pmod{n}$. Suppose that $k \in [n]$. For each $i \in [n]$, we have $g^k(i) = (i + k) \bmod n$. In particular, $g^n(i) = i$, so that $g^n = i_{[n]} = (1)(2) \cdots (n)$. Note first that for every $k \in [n]$ there exists a divisor d of n such that $\gcd(k, n) = d$, and for every divisor d of n there exists a $k \in [n]$ such that $\gcd(k, n) = d$. Partition the members of $[n]$ into the classes $C_d := \{k \in [n] : \gcd(k, n) = d\}$. By Theorem 8.1.4, $|C_d| = \phi(n/d)$. Suppose now that $k \in C_d$. Then g^k consists of d cycles of length n/d, by the following argument. If $\gcd(k, n) = d$, then $k = jd$, for some j such that $\gcd(j, n/d) = 1$. Each cycle of g^k takes the form

$$(i, (i + jd) \bmod n, (i + 2jd) \bmod n, \ldots, (i + (\frac{n}{d} - 1)jd) \bmod n), \tag{10.2.6}$$

since $(i + \frac{n}{d} jd) \bmod n = i$ (so the length of this cycle is at most n/d) . In fact, the length of this cycle is exactly n/d, since members of the above cycle are distinct. For if $i + rjd \equiv i + sjd (\mathrm{mod} n)$, for $0 \le r < s \le \frac{n}{d} - 1$, then $n = \frac{n}{d} d | (r - s)jd$, and so $\frac{n}{d} | (r - s)j$. Since j and $\frac{n}{d}$ are relatively prime, we must have $\frac{n}{d} | (r - s)$. But this is impossible, since $0 < (r - s) < \frac{n}{d}$. But the cycles of g^k correspond to orbits of the cyclic group generated by g^k, and since such orbits partition the set $[n]$, the distinct cycles of g^k are disjoint and exhaust $[n]$. So g^k has d cycles, each of length n/d. Formula (10.2.5) follows immediately from the definition of the cycle index. □

Note that formula (8.2.2), which you were asked to derive from formula (8.2.1) in Exercise 8.1, is an immediate consequence of Theorem 10.2.4.

(iv) Suppose that $n \geq 3$. The *dihedral group* D_n consists of the $2n$ symmetries (rotations and reflections) of a regular n-gon, with sides labeled $1, 2, \ldots, n$. D_n contains n rotational symmetries, corresponding to clockwise rotations through the angles $2\pi k/n$, for $0 \leq k \leq n-1$. If n is odd, each axis of symmetry (across which a reflection takes place) connects the midpoint of a side to the opposite vertex. If n is even, there are $n/2$ axes of symmetry connecting the midpoints of opposite sides, and $n/2$ axes of symmetry connecting opposite vertices.

Theorem 10.2.5. *Suppose that $n \geq 3$. If n is odd, then*

$$Z(D_n; z_1, \ldots, z_n) = \frac{1}{2} Z(C_n; z_1, \ldots, z_n) + \frac{1}{2} z_1 z_2^{(n-1)/2}. \tag{10.2.7}$$

If n is even, then

$$Z(D_n; z_1, \ldots, z_n) = \frac{1}{2} Z(C_n; z_1, \ldots, z_n) + \frac{1}{4} (z_1^2 z_2^{(n-2)/2} + z_2^{n/2}). \tag{10.2.8}$$

Proof. Exercise. □

As an application of the above theorem, consider the following *necklace counting problem*. In how many ways can you construct a necklace using n beads, each of a different color? It is assumed that the beads are strung together with no noticeable clasp, and that they are uniform in color, with no decoration that must be face up. Clearly, any two of the $n!$ visibly distinct ways of arranging the beads around, say, the sides of a regular n-gon, give rise to the (for all practical purposes) same necklace if one arrangement results from any of the n possible rotations, or n possible reflections, of the other. So the classes of visibly distinct arrangements that give rise to the same necklace have uniform cardinality $2n$, and the number of necklaces is $n!/2n$. Suppose, however, that you must construct a necklace consisting of n beads, that there are at least n beads of each of k different colors available, and that there are no constraints on the colors to be used. In particular, necklaces with all beads of the same color are allowed. Here, unlike the case just considered, different rotations or reflections of an arrangement of beads around the sides of an n-gon need not produce an arrangement visually distinct from the former. So the full force of Theorem 10.2.2, along with formulas (10.2.7) and (10.2.8), needs to be applied to necklace counting, as you are asked to do in one of the exercises.

10.3. The pattern inventory: Pólya's second theorem

In what follows, we simplify notation by taking $D = [n]$ and $C = [k]$. If $\varphi : [n] \to [k]$, the *frequency type* of φ is the sequence

$$(f_1, \ldots, f_k) = (|\varphi^{-1}(\{1\})|, \ldots, |\varphi^{-1}(\{k\}|). \tag{10.3.1}$$

If $\varphi_1 \sim \varphi_2$ relative to some permutation group G on $[n]$, the colorings φ_1 and φ_2 clearly have the same frequency type (you are asked to prove this in Exercise 10.1(a)), and so we may speak of the frequency type of each of the G-classes induced by $\sim$. In this

section we investigate the *pattern inventory for G-classes*, i.e., the polynomial

(10.3.2)
$$P(G; c_1, \dots, c_k) := \sum_{\substack{f_1+\cdots+f_k=n \\ f_j \geq 0}} [\#\text{ of } G\text{-classes with frequency type } (f_1, \dots, f_k)] c_1^{f_1} \cdots c_k^{f_k},$$

with $c_1, \dots, c_k$ construed as indeterminates. In order to determine the pattern inventory, we need a weighted version of Burnside's lemma.

Let G be a permutation group on the finite set X. A map $w : X \to A$, where A is a commutative ring containing $\mathbb{Q}$, is called a *weight function on X with respect to G* if w is constant on the orbits of G. A weight function w on X induces a map $W : O(G) \to A$ in the obvious way.

Theorem 10.3.1 (Weighted Burnside's lemma)**.** *In the above context,*

$$\sum_{O \in O(G)} W(O) = \frac{1}{|G|} \sum_{g \in G} \sum_{\substack{x \in X \\ g(x)=x}} w(x). \tag{10.3.3}$$

Proof.

$$\begin{aligned} \sum_{g \in G} \sum_{\substack{x \in X \\ g(x)=x}} w(x) &= \sum_{x \in X} w(x) \sum_{\substack{g \in G \\ g(x)=x}} 1 = \sum_{x \in X} w(x) |G_x| \\ &= \sum_{x \in X} w(x) \frac{|G|}{|O_x|} = |G| \sum_{O \in O(G)} \sum_{x \in O} \frac{w(x)}{|O|} \\ &= |G| \sum_{O \in O(G)} W(O), \end{aligned}$$

since $\sum_{x \in O} \frac{w(x)}{|O|}$ consists of $|O|$ terms, each equal to $\frac{W(O)}{|O|}$. □

We are now in the position to determine the pattern inventory $P(G; c_1, \dots, c_k)$.

Theorem 10.3.2 (Pólya's second theorem)**.**

$$P(G; c_1, \dots, c_k) = Z(G; \sum_{j=1}^{k} c_j, \sum_{j=1}^{k} c_j^2, \dots, \sum_{j=1}^{k} c_j^n). \tag{10.3.4}$$

Proof. For each coloring φ, let

$$w(\varphi) := \prod_{i=1}^{n} c_{\varphi(i)} = c_1^{f_1(\varphi)} c_2^{f_2(\varphi)} \cdots c_k^{f_k(\varphi)}, \tag{10.3.5}$$

where $f_j(\varphi) = |\varphi^{-1}(\{j\})|$, $j = 1, \dots, k$. As usual, G-classes of colorings are identical with orbits of $\overline{G}$, the permutation group on $[k]^{[n]}$ induced by G. Moreover, w, as defined by (10.3.5), is constant on the orbits of $\overline{G}$. If W is the weight function on orbits derived

from w, then, clearly, $P(G; c_1, \ldots, c_k) = \sum_{O \in O(\overline{G})} W(O)$. But, by Theorem 10.3.1,

$$\sum_{O \in O(\overline{G})} W(O) = \frac{1}{|G|} \sum_{g \in G} \sum_{\substack{\varphi \in [k]^{[n]} \\ \overline{g}(\varphi) = \varphi}} w(\varphi) = \frac{1}{|G|} \sum_{g \in G} \sum_{\substack{\varphi \in [k]^{[n]} \\ \varphi \circ g = \varphi}} \prod_{i=1}^{n} c_{\varphi(i)}$$

$$\frac{1}{|G|} \sum_{g \in G} (c_1 + \cdots + c_k)^{\lambda_1(g)} (c_1^2 + \cdots + c_k^2)^{\lambda_2(g)} \cdots (c_1^n + \cdots + c_k^n)^{\lambda_n(g)}$$

$$= Z(G; \sum_{j=1}^{k} c_j, \sum_{j=1}^{k} c_j^2, \ldots, \sum_{j=1}^{k} c_j^n),$$

since, as previously noted, $\varphi \circ g = \varphi$ if and only if φ is constant on the cycles of g. □

Remark 10.3.3. If we set $c_1 = \cdots = c_k = 1$, then every coloring, hence every orbit of $\overline{G}$, receives weight 1, and (10.3.4) reduces to Pólya's first theorem. What is the effect of setting a proper subset of the aforementioned indeterminates equal to 1 and leaving the others unchanged? What is the effect of setting a proper subset of these indeterminates equal to 0?

10.4. Counting isomorphism classes of graphs

We have previously defined a *graph* on a set V to be an irreflexive, symmetric binary relation on V. In this section, it will be convenient to use an equivalent conceptualization of a such a graph as an ordered pair (V, E), where $E \subseteq \binom{V}{2}$, the set of all two-element subsets of V. Elements of V are called *vertices*, and elements of E are called *edges*. The graphs (V_1, E_1) and (V_2, E_2) are *isomorphic* (symbolized by $(V_1, E_1) \cong (V_2, E_2)$) if there is a bijection $\sigma : V_1 \to V_2$ such that $\{x, y\} \in E_1 \Leftrightarrow \{\sigma(x), \sigma(y)\} \in E_2$. Henceforth, we shall take $V = [n]$, where $n \geq 3$, and attempt to enumerate the isomorphism classes of graphs on $[n]$. This is often referred to as the problem of counting *graphs on n unlabeled vertices*, or enumerating *nonisomorphic graphs on* $[n]$. To apply Pólya theory to this problem, we view each graph $([n], E)$ as a coloring $\varphi : \binom{[n]}{2} \to \{yes, no\}$, where $\varphi(\{x, y\}) = yes$ if $\{x, y\} \in E$ and $\varphi(\{x, y\}) = no$ if $\{x, y\} \notin E$. As usual, S_n denotes the symmetric group on $[n]$, and $S(\binom{[n]}{2})$ denotes the symmetric group on $\binom{[n]}{2}$. If $\sigma \in S_n$ and $\{x, y\} \in \binom{[n]}{2}$, let

$$g_\sigma(\{x, y\}) := \{\sigma(x), \sigma(y)\}. \tag{10.4.1}$$

Since σ is injective, $g_\sigma : \binom{[n]}{2} \to \binom{[n]}{2}$. Also, it is easy to see that g_σ is injective, and hence bijective. So the map $\sigma \mapsto g_\sigma$ is a map from S_n into $S(\binom{[n]}{2})$, and, indeed, a homomorphism, i.e., $g_{\sigma_1 \circ \sigma_2} = g_{\sigma_1} \circ g_{\sigma_2}$. In fact, since we are assuming that $n \geq 3$, this map is injective. So the range of this map, which we denote by $S_n^{(2)}$, called the *pair group*, is actually isomorphic to S_n.

Now we have captured the notion of an isomorphism class of graphs in a Pólya framework: Given graphs (construed as colorings) $\varphi_1, \varphi_2 : \binom{[n]}{2} \to \{yes, no\}$, $\varphi_1 \cong \varphi_2 \Leftrightarrow \exists \sigma \in S_n$ such that, for all $\{x, y\} \in \binom{[n]}{2}$,

$$\varphi_1(\{x, y\}) = \varphi_2(\{\sigma(x), \sigma(y)\}) = \varphi_2(g_\sigma(\{x, y\})) \Leftrightarrow \varphi_1 = \varphi_2 \circ g_\sigma \Leftrightarrow \varphi_1 \sim \varphi_2.$$

So isomorphism classes of graphs are just G-classes of colorings with respect to the permutation group $G = S_2^{(n)}$. It follows from Pólya's first theorem that the number of isomorphism classes of graphs on $[n]$ is given by

$$Z(S_2^{(n)}; z_1, \ldots, z_{\binom{n}{2}})|_{z_i=2}. \tag{10.4.2}$$

Furthermore, by Pólya's second theorem (representing the "color" *yes* by the indeterminate e (for *edge*), and substituting 1 for whatever indeterminate is chosen to represent *no*, we have

$$\begin{aligned}\sum_{j=0}^{\binom{n}{2}} &(\#\text{of isomorphism classes of graphs on } [n] \text{with exactly j edges}) e^j \\ &= Z(S_2^{(n)}; 1+e, 1+e^2, \ldots, 1+e^{\binom{n}{2}}).\end{aligned} \tag{10.4.3}$$

So we need only determine the cycle index of the permutation group $S_n^{(2)}$ in order to enumerate the isomorphism classes of graphs on $[n]$, and the isomorphism classes of such graphs with a fixed number of edges. Since S_n and $S_n^{(2)}$ are isomorphic as groups, one might hastily suppose that this problem has already been solved. But the structures of S_n and $S_n^{(2)}$ as *permutation groups* are quite distinct, and their cycle indices rarely coincide. For example, while

$$Z(S_3; z_1, z_2, z_3) = Z(S_3^{(2)}; z_1, z_2, z_3) = \frac{1}{6}(z_1^3 + 3z_1z_2 + 2z_3), \tag{10.4.4}$$

it is the case that

$$Z(S_4; z_1, \ldots, z_4) = \frac{1}{24}(z_1^4 + 6z_1^2z_2 + 8z_1z_3 + 3z_2^2 + 6z_4), \tag{10.4.5}$$

whereas

$$Z(S_4^{(2)}; z_1, \ldots, z_6) = \frac{1}{24}(z_1^6 + 6z_1^2z_2^2 + 8z_3^2 + 3z_1^2z_2^2 + 6z_2z_4). \tag{10.4.6}$$

We have avoided combining the terms $6z_1^2z_2^2$ and $3z_1^2z_2^2$ in (10.4.6) since doing so would obscure the correspondence between the cycle structures of S_n and $S_n^{(2)}$. Suppose that $\sigma \in S_n$ and $g_\sigma \in S_m^{(2)}$, as given by (10.4.1) above. It turns out that cycles of g_σ arise from cycles of σ in one of two ways, *intracyclically* (i.e., within a cycle), and *intercyclically* (between two cycles). Before treating the general case, we consider a few concrete numerical examples.

(1) Suppose that σ contains a cycle of odd length, say (1 2 3 4 5). Then the following are cycles of g_σ:
   (i) $(\{1,2\}\{2,3\}\{3,4\}\{4,5\}\{5,1\})$ and
   (ii) $(\{1,3\}\{2,4\}\{3,5\}\{4,1\}\{5,2\})$.
   These are the only cycles of g_σ whose members are pairs of vertices, both of which are members of the set $\{1,2,3,4,5\}$. For there are $\binom{5}{2} = 10$ such pairs, and each appears exactly once in one of the above two cycles.

(2) Suppose that σ contains a cycle of even length, say (1 2 3 4 5 6). Then the following are cycles of g_σ:
   (i) $(\{1,2\}\{2,3\}\{3,4\}\{4,5\}\{5,6\}\{6,1\})$,

(ii) $(\{1,3\}\{2,4\}\,\{3,5\}\{4,6\}\{5,1\}\{6,2\})$, and
(iii) $(\{1,4\}\{2,5\}\{3,6\})$.
These are the only cycles of g_σ whose members are pairs of vertices, both of which are members of the set $\{1,2,3,4,5,6\}$. For there are $\binom{6}{2} = 15$ such pairs, and each appears exactly once in one of the above three cycles. The foregoing examples illustrate the intracyclic generation of cycles of g_σ from cycles of σ. The next two examples illustrate the intercyclic generation of such cycles.

(3) Suppose that σ contains the cycle (1 2 3 4 5) as well as the cycle (6). Then the following is a cycle of g_σ: $(\{1,6\}\{2,6\}\{3,6\}\{4,6\}\{5,6\})$. This is the only cycle of g_σ whose members are pairs of vertices, one of which comes from the set $\{1,2,3,4,5\}$ and the other from the set $\{6\}$. For there are $5 \cdot 1 = 5$ such pairs, and each appears exactly once in the above cycle.

(4) Suppose that σ contains the cycles (1 2 3 4) and (5 6 7 8 9 10). Then the following are cycles of g_σ:
(i) $(\{1,5\}\{2,6\}\{3,7\}\{4,8\}\{1,9\}\{2,10\}\{3,5\}\{4,6\}\{1,7\}\,\{2,8\}\{3,9\}\{4,10\})$ and
(ii) $(\{1,6\}\{2,7\}\{3,8\}\{4,9\}\{1,10\}\{2,5\}\{3,6\}\{4,7\}\{1,8\}\{2,9\}\{3,10\}\{4,5\})$.
These are the only cycles of g_σ whose members are pairs of vertices, one of which comes from the set $\{1,2,3,4\}$ and the other from the set $\{5,6,7,8,9,10\}$. For there are $4 \cdot 6 = 24$ such pairs, and each appears exactly once in one the above two cycles.

We are now prepared to treat the general case. First, recall that if $m \in \mathbb{P}$ and $a, b \in \mathbb{Z}$, the relation $a \equiv b \pmod{m}$ holds if and only if $m|b-a$. Varying our notation slightly from that in section 8.2, we denote by $a \bmod m$ the unique $r \in [m]$ (*not* in $\{0,1,\ldots,m-1\}$) such that $a \equiv r \pmod{m}$. In what follows, $\sigma \in S_n$ and $g_\sigma \in S_n^{(2)}$ is defined by (10.4.1). Lemmas 10.4.1–10.4.3 describe how cycles of g_σ arise from cycles of σ.

Lemma 10.4.1. *If $m \in \mathbb{N}$, each cycle $(v_1 v_2 \cdots v_{2m+1})$ of σ gives rise in g_σ to m cycles of length $2m+1$, where the members of the latter cycles are pairs of vertices, each belonging to $\{v_1, v_2, \ldots, v_{2m+1}\}$.*

Proof. Consider an arbitrary pair $\{v_j, v_k\}$, with j, k distinct members of $[2m+1]$. We claim that the cycle of g_σ containing $\{v_j, v_k\}$ is

$$(\{v_{j+i-1 \bmod 2m+1}, v_{k+i-1 \bmod 2m+1}\})_{1 \le i \le 2m+1}. \tag{10.4.7}$$

That $g_\sigma(\{v_{j+i-1 \bmod 2m+1}, v_{k+i-1 \bmod 2m+1}\}) = \{v_{j+i \bmod 2m+1}, v_{k+i \bmod 2m+1}\}$ for $1 \le i \le 2m+1$ is clear. To show that (10.4.7) defines a cycle of length $2m+1$, we need to show that if $i, i' \in [2m+1]$ and

$$\{v_{j+i-1 \bmod 2m+1}, v_{k+i-1 \bmod 2m+1}\} = \{v_{j+i'-1 \bmod 2m+1}, v_{k+i'-1 \bmod 2m+1}\}, \tag{10.4.8}$$

then $i = i'$. If (10.4.8) holds, this cannot be in virtue of the fact that $v_{j+i-1 \bmod 2m+1} = v_{k+i'-1 \bmod 2m+1}$ and $v_{k+i-1 \bmod 2m+1} = v_{j+i'-1 \bmod 2m+1}$, for this would imply that $j+i-1 \equiv k+i'-1 \pmod{2m+1}$ and $k+i-1 \equiv j+i'-1 \pmod{2m+1}$, whence $j \equiv k \pmod{2m+1}$, and so $j = k$, a contradiction. So we must have, *inter alia*, that $v_{j+i-1 \bmod 2m+1} = v_{j+i'-1 \bmod 2m+1}$, whence $j+i-1 \equiv j+i'-1 \pmod{2m+1}$ and

thus $i \equiv i' \pmod{2m+1}$, and so $i = i'$. The cycles of g_σ partition the $\binom{2m+1}{2}$-element set $\binom{\{v_1, v_2, \ldots, v_{2m+1}\}}{2}$ into blocks of cardinality $2m+1$ in the obvious way.

So there are $\binom{2m+1}{2}/(2m+1) = m$ such cycles. □

Lemma 10.4.2. *If $m \in \mathbb{P}$, a cycle $(v_1 v_2 \cdots v_{2m})$ of σ gives rise in g_σ to $m-1$ cycles of length $2m$ and one cycle of length m, where the members of the latter cycles are pairs of vertices, each belonging to the set $\{v_1, v_2, \ldots, v_{2m}\}$.*

Proof. Consider again an arbitrary pair $\{v_j, v_k\}$, with j, k distinct members of $[2m]$. If $j \equiv k \pmod{m}$, then the cycle containing $\{v_j, v_k\}$ is

$$(\{v_1, v_{m+1}\}\{v_2, v_{m+2}\} \cdots \{v_m, v_{2m}\}), \tag{10.4.9}$$

i.e., all such pairs belong to a single cycle of length m. If j and k are not congruent $\pmod{m}$, then the cycle containing $\{v_j, v_k\}$ is

$$(\{v_{j+i-1 \bmod 2m}, v_{k+i-1 \bmod 2m}\})_{1 \le i \le 2m}. \tag{10.4.10}$$

By an argument nearly identical to that in the proof of Lemma 10.4.1, it may be shown that the pairs appearing in (10.4.10) are distinct, and so (10.4.10) is indeed a cycle of g_σ of length $2m$. The distinct cycles of this type, along with the cycle (10.4.9), partition the $\binom{2m}{2}$-element set $\binom{\{v_1, \ldots, v_{2m}\}}{2}$ of pairs of vertices belonging to $\{v_1, \ldots, v_{2m}\}$ into one cycle of cardinality m and x cycles of cardinality $2m$, whence $m + 2mx = \binom{2m}{2}$, and so $x = m-1$. □

Lemma 10.4.3. *If $r, s \in \mathbb{P}$, each pair $\{(u_1 \cdots u_r), (v_1 \cdots v_s)\}$ of cycles of σ gives rise in g_σ to $\gcd(r,s)$ cycles of length $\operatorname{lcm}(r,s)$, where members of the latter cycle are pairs of vertices $\{u_j, v_k\}$, with u_j belonging to the cycle $(u_1 \cdots u_r)$ and v_k to the cycle $(v_1 \cdots v_s)$.*

Proof. We claim that the cycle containing the pair $\{u_j, v_k\}$ is

$$(\{u_{j+i-1 \bmod r}, v_{k+i-1 \bmod s}\})_{1 \le i \le \operatorname{lcm}(r,s)}. \tag{10.4.11}$$

To demonstrate this, we need to show that

$$g_\sigma(\{u_{j+i-1 \bmod r}, v_{k+i-1 \bmod s}\}) = \{u_{j+i \bmod r}, v_{k+i \bmod s}\}$$

for $1 \le i \le \operatorname{lcm}(r,s)$, and that the pairs in (10.4.11) are distinct. The proof of this fact is left as an exercise. Since the rs pairs of vertices, one from $\{u_1, \ldots, u_r\}$ and the other from $\{v_1, \ldots, v_s\}$ are partitioned into cycles of length $\operatorname{lcm}(r,s)$, there must be $rs/\operatorname{lcm}(r,s) = \gcd(r,s)$ such cycles. □

The above propositions yield the following algorithm for constructing the cycle index of $S_n^{(2)}$ from the cycle index of S_n.

Theorem 10.4.4.

(i) *A factor z_{2m+1} in a term of the cycle index of S_n gives rise to the factor z_{2m+1}^m in the corresponding term of the cycle index of $S_n^{(2)}$.*

(ii) *A factor z_{2m} in a term of the cycle index of S_n gives rise to the factor $z_m z_{2m}^{m-1}$ in the corresponding term of the cycle index of $S_n^{(2)}$.*

(iii) *A factor $z_r z_s$ in a term of the cycle index of S_n gives rise to the factor $z_{\mathrm{lcm}(r,s)}^{\gcd(r,s)}$ in the corresponding term of the cycle index of $S_n^{(2)}$.*

Proof. Obvious. □

To illustrate the application of Theorem 10.4.4, we derive the cycle index of $S_3^{(2)}$. We start with the fact that $Z(S_3; z_1, z_2, z_3) = \frac{1}{6}(z_1^3 + 3z_1z_2 + 2z_3)$.

(*i*) each of the three factors z_1 in the term z_1^3 gives rise intracyclically to the factor $z^0 = 1$ in the corresponding term of the cycle index of $S_3^{(2)}$, and each of the three pairs of factors $\{z_1, z_1\}$ in the term z_1^3 gives rise intercyclically to the factor $(z_1)_{\mathrm{lcm}(1,1)}^{\gcd(1,1)} = z_1$. So the term z_1^3 in the cycle index of S_3 gives rise to the corresponding term z_1^3 in the cycle index of $S_3^{(2)}$.

(*ii*) The factor z_1 in the term $3z_1z_2$ gives rise intracyclically to the factor 1, and the term z_2 to the factor z_1. The pair of factors $\{z_1, z_2\}$ give rise intercyclically to the term $z_{\mathrm{lcm}(1,2)}^{\gcd(1,2)} = z_2$. So the term $3z_1z_2$ in the cycle index of S_3 gives rise to the corresponding term $3z_1z_2$ in the cycle index of $S^{(2)}$.

(*iii*) The factor z_3 in the term $2z_3$ gives rise intracyclically to the factor $z_3^1 = z_3$. So the term $2z_3$ in the cycle index of S_3 gives rise to the corresponding term $2z_3$ in the cycle index of $S_3^{(2)}$.

In Exercise 10.9 you are asked to derive the cycle index of $S_4^{(2)}$. The cycle indices of $S_n^{(2)}$ have been tabulated for many values of n; see, for example, Harary and Palmer (1973).

10.5. G-classes of proper subsets of colorings / group actions

Suppose now that Φ is a *proper subset* of C^D, G is a permutation group on D, and we wish to determine the number of G-classes into which the equivalence relation $\sim$ partitions Φ. In such a case, the permutation group $\overline{G}$ induced on Φ by G may differ from G in its cardinality. Nevertheless, there is a formula (though not as elegant as the formula given by Pólya's first theorem) which can be used to enumerate such G-classes. Its proof requires the following extension of Burnside's lemma.

Theorem 10.5.1 (Burnside's lemma—extended)**.** *Let X be a finite nonempty set, and let G be any finite group (not necessarily a permutation group on X). Suppose that the map $g \mapsto \overline{g}$ is a homomorphism from G into the symmetric group $S(X)$, from which it follows that the range $\overline{G}$ of this map is a subgroup of $S(X)$. Then*

$$|O(\overline{G})| = \frac{1}{|G|} \sum_{g \in G} \lambda_1(\overline{g}). \tag{10.5.1}$$

Proof. Let K be the kernel of the map $g \mapsto \overline{g}$, with $|K| = m$. Then each $\pi \in \overline{G}$ has m preimages in G under this map, whence $|G| = m|\overline{G}|$ and $\sum_{g \in G} \lambda_1(\overline{g}) = m \sum_{\sigma \in \overline{G}} \lambda_1(\sigma)$.

By Burnside's Lemma,

$$|O(\overline{G})| = \frac{1}{|\overline{G}|}\sum_{\sigma\in\overline{G}}\lambda_1(\sigma) = \frac{m}{|G|}\frac{1}{m}\sum_{g\in G}\lambda_1(\overline{g}) = \frac{1}{|G|}\sum_{g\in G}\lambda_1(\overline{g}). \qquad \square$$

Remark. When, as above, there is a homomorphism from the group G into $S(X)$, we say that G *acts on* X.

Suppose now that $\Phi \subseteq C^D$, that G is a permutation group on D, and that for all $\varphi \in \Phi$ and all $g \in G$, $\varphi \circ g \in \Phi$. As before, if $\varphi_1, \varphi_2 \in \Phi$, we write $\varphi_1 \sim \varphi_2$ if there exists a $g \in G$ such that $\varphi_1 \circ g = \varphi_2$. To invoke Theorem 10.5.1 in the enumeration of the G-classes into which Φ is partitioned by $\sim$, we define a new group operation $*$ on G by $g_1 * g_2 := g_2 \circ g_1$. Then the map $g \mapsto \overline{g}$ from the group $(G, *)$ into $S(\Phi)$, where, as before, $\overline{g}(\varphi) = \varphi \circ g$, is a homomorphism, with range denoted by $\overline{G}$. And again, G-classes are identical with orbits of $\overline{G}$. This leads to the following extension of Pólya's first theorem.

Theorem 10.5.2 (Pólya's first theorem—extended)**.** *The number of G-classes into which Φ is partitioned by $\sim$ is equal to*

$$\frac{1}{|G|}\sum_{g\in G}|\{\varphi\in\Phi : \varphi \text{ is constant on the cycles of } g\}| \tag{10.5.2}$$

$$= \frac{1}{|G|}\sum_{\lambda_1+2\lambda_2+\cdots+n\lambda_n=n}|\{g\in G : \lambda_i(g)=\lambda_i, i=1,\dots,n\}| \times |\{\varphi\in\Phi : \varphi \text{ is constant on the cycles of a permutation of type } 1^{\lambda_1}2^{\lambda_2}\cdots n^{\lambda_n}\}|.$$

Proof. Exercise. $\square$

Remark 10.5.3. When $\Phi = C^D$, formula (10.5.2) reduces to $Z(G; k, \dots, k)$. In general, however, (10.5.2) must be evaluated on a case-by-case basis, depending on the nature of Φ and G. Exercise 10.12 features a simple application of Theorem 10.5.2.

10.6. De Bruijn's generalization of Pólya theory

Suppose that G is a permutation group on D and that H is a permutation group on C. If $\varphi_1, \varphi_2 : D \to C$, we write $\varphi_1 \approx \varphi_2$ if $\exists g \in G$ and $\exists h \in H$ such that $\varphi_2 = h \circ \varphi_1 \circ g$. The relation $\approx$ is clearly an equivalence relation. Our goal is to determine the number of equivalence classes into which $\approx$ partitions the set of colorings C^D. The solution to this problem features another application of Theorem 10.5.1.

We first define a binary composition $*$ on the set $G \times H$ by

$$(g_1, h_1) * (g_2, h_2) := (g_2 \circ g_1, h_1 \circ h_2). \tag{10.6.1}$$

It is easy to check that $(G \times H, *)$ is a group. If $(g, h) \in G \times H$ and $\varphi \in C^D$, define $\overline{(g,h)} : C^D \to C^D$ by

$$\overline{(g,h)}(\varphi) := h \circ \varphi \circ g. \tag{10.6.2}$$

The map $\overline{(g,h)}$ is clearly injective and, hence, bijective, i.e., a permutation on C^D. Moreover, the map $(g,h) \mapsto \overline{(g,h)}$ is a homomorphism from $(G\times H, *)$ into $(S(C^D), \circ)$, since, for each $\varphi \in C^D$,

$$\begin{aligned}\overline{(g_1,h_1)*(g_2,h_2)}(\varphi) &= \overline{(g_2\circ g_1, h_1\circ h_2)}(\varphi)\\ &= (h_1\circ h_2)\circ\varphi\circ(g_2\circ g_1) = h_1\circ(h_2\circ\varphi\circ g_2)\circ g_1\\ &= h_1\circ(\overline{(g_2,h_2)}(\varphi))\circ g_1 = \overline{(g_1,h_1)}\circ\overline{(g_2,h_2)}(\varphi).\end{aligned}$$

Denote by $\overline{G\times H}$ the range of the homomorphism $(g,h)\mapsto\overline{(g,h)}$, whence $\overline{G\times H}$ is a subgroup of $S(C^D)$, i.e., a permutation group on C^D. Clearly, $\varphi_1 \approx \varphi_2$ if and only if $\varphi_1 \equiv \varphi_2 (\mathrm{mod}\,\overline{G\times H})$.

Lemma 10.6.1. *The number of equivalence classes into which $\approx$ partitions the set of colorings C^D is equal to*

$$\frac{1}{|G|}\frac{1}{|H|}\sum_{(g,h)\in G\times H}\ \sum_{\substack{\varphi\in C^D\\ h\circ\varphi\circ g=\varphi}} 1. \tag{10.6.3}$$

Proof. By the above remarks, along with Theorem 10.5.1, it follows that the number of $\approx$ equivalence classes is equal to

$$\begin{aligned}|O_{\overline{G\times H}}| &= \frac{1}{|G\times H|}\sum_{(g,h)\in G\times H}\lambda_1(\overline{(g,h)})\\ &= \frac{1}{|G|}\frac{1}{|H|}\sum_{(g,h)\in G\times H}\ \sum_{\substack{\varphi\in C^D\\ \overline{(g,h)}(\varphi)=\varphi}} 1 = \frac{1}{|G|}\frac{1}{|H|}\sum_{(g,h)\in G\times H}\ \sum_{\substack{\varphi\in C^D\\ h\circ\varphi\circ g=\varphi}} 1.\end{aligned}$$

□

Lemma 10.6.2. *If $(g,h)\in G\times H$, a coloring $\varphi : [n]\to[k]$ satisfies $h\circ\varphi\circ g = \varphi$ if and only if φ maps the elements of any cycle of g of length i onto elements of a cycle of h of length j, where $j|i$, and in such a way that if $\varphi(b) = c$, then $\varphi(g(b)) = h^{-1}(c)$.*

Proof. *Sufficiency*. Let $b\in[n]$, and suppose that $\varphi(b) = c$. Then $\varphi(g(b)) = h^{-1}(c)$, and so $h(\varphi(g(b))) = c = \varphi(b)$. *Necessity*. Let $(b, g(b), g^2(b), \ldots, g^{i-1}(b))$ be a cycle of g of length i, and suppose that $\varphi(b) = c$. Since $\varphi(x) = h(\varphi(g(x)))$ for all $x\in[n]$, $\varphi(g(x)) = h^{-1}(\varphi(x))$ for all $x\in[n]$. So $(\varphi(b), \varphi(g(b)), \varphi(g^2(b)), \ldots, \varphi(g^{i-1}(b))) = (c, h^{-1}(c), (h^{-1})^2(c), \ldots, (h^{-1})^{i-1}(c))$. Since $g^i(b) = b$, $\varphi(g^i(b)) = \varphi(b) = c$, and so $(h^{-1})^i(c) = c$. So while the elements of $(c, h^{-1}(c), (h^{-1})^2(c), \ldots, (h^{-1})^{i-1}(c))$ need not be distinct, they do exhaust all the elements of a cycle of h^{-1}, and hence of h. And the length of that cycle, call it j, is obviously a divisor of i. □

Theorem 10.6.3 (De Bruijn's theorem)**.** *The number of equivalence classes into which* $\approx$ *partitions* C^D *is given by the formula*

$$\frac{1}{|G|}\frac{1}{|H|}\sum_{(g,h)\in G\times H}\prod_{i=1}^{n}\Big(\sum_{j|i} j\lambda_j(h)\Big)^{\lambda_i(g)} \tag{10.6.4}$$

$$\begin{aligned}&= Z\Big(G;\frac{\partial}{\partial z_1},\frac{\partial}{\partial z_2},\dots,\frac{\partial}{\partial z_n}\Big)\\ &\quad\times Z\big(H;e^{z_1+z_2+z_3+\cdots},e^{2(z_2+z_4+z_6+\cdots)},\\ &\qquad e^{3(z_3+z_6+z_9+\cdots)},\dots,e^{k(z_k+z_{2k}+z_{3k}+\cdots)}\big)\big|_{z_1=z_2=\cdots=z_n=0}.\end{aligned} \tag{10.6.5}$$

Proof. (De Bruijn (1959)). Formula (10.6.4) is an immediate consequence of Lemmas 10.6.1 and 10.6.2. Given each cycle of g of length i (henceforth, i-cycle), a map of its elements to any j-cycle of h is completely determined by choosing any d in the i-cycle and setting $\varphi(d)$ equal to any element of the j-cycle in question. To derive (10.6.5), it is convenient to abbreviate $\lambda_i(g)$, the number of i-cycles of g, by g_i, and $\lambda_j(h)$ by h_j. Then (10.6.4) becomes

$$\begin{aligned}&\frac{1}{|G|}\frac{1}{|H|}\sum_{(g,h)\in G\times H}\prod_{i=1}^{n}\Big(\sum_{j|i} jh_j\Big)^{g_i}\\ &\quad=\frac{1}{|G|}\frac{1}{|H|}\sum_{(g,h)\in G\times H}(h_1)^{g_1}(h_1+2h_2)^{g_2}(h_1+3h_3)^{g_3}\\ &\qquad\qquad\times(h_1+2h_2+4h_4)^{g_4}\cdots(h_1+\cdots+nh_n)^{g_n}.\end{aligned}$$

But,

(i) $(h_1)^{g_1}=\Big(\frac{\partial}{\partial z_1}\Big)^{g_1}e^{h_1z_1}\big|_{z_1=0}$,

(ii) $(h_1+2h_2)^{g_2}=\Big(\frac{\partial}{\partial z_2}\Big)^{g_2}e^{(h_1+2h_2)z_2}\big|_{z_2=0}=\Big(\frac{\partial}{\partial z_2}\Big)^{g_2}e^{h_1z_2}e^{2h_2z_2}\big|_{z_2=0}$,

(iii) $(h_1+3h_3)^{g_3}=\Big(\frac{\partial}{\partial z_3}\Big)^{g_3}e^{h_1z_3}e^{3h_3z_3}\big|_{z_3=0}$, etc.,

which yields (10.6.5). □

10.7. Equivalence classes of boolean functions

A function $\varphi:\{0,1\}^n\to\{0,1\}$ is called an *n-ary boolean function* (or *n-ary truth function*, if 1 is regarded as designating "true" and 0 as designating "false"). Examples, which will be familiar to those who have studied propositional logic, include the following:

- the *negation function* $\varphi(p_1)=\neg p_1$, where $\neg 1=0$ and $\neg 0=1$;
- the *disjunction function* $\varphi(p_1,p_2)=p_1\vee p_2$, where $1\vee 1=1\vee 0=0\vee 1=1$ and $0\vee 0=0$; and
- the *conjunction function* $\varphi(p_1,p_2)=p_1\wedge p_2$, where $1\wedge 1=1$ and $1\wedge 0=0\wedge 1=0\wedge 0=0$.

Every n-ary boolean function can be expressed using negation, disjunction, and conjunction, as illustrated by the following example. Suppose that $\varphi(0,0,1) = \varphi(0,1,0) = \varphi(0,1,1) = \varphi(1,0,1) = \varphi(1,1,0) = \varphi(1,1,1) = 1$, with $\varphi(p_1,p_2,p_3) = 0$ in all other cases. It is easy to verify that

$$\varphi(p_1,p_2,p_3) = (\neg p_1 \wedge \neg p_2 \wedge p_3) \vee (\neg p_1 \wedge p_2 \wedge \neg p_3) \vee (\neg p_1 \wedge p_2 \wedge p_3) \\ \vee (p_1 \wedge \neg p_2 \wedge p_3) \vee (p_1 \wedge p_2 \wedge \neg p_3) \vee (p_1 \wedge p_2 \wedge p_3),$$

which is called the *disjunctive normal form of* φ. It should be clear on the basis of the preceding example how to construct the disjunctive normal form of any n-ary boolean function from the n-tuples in the preimage $\varphi^{-1}(\{1\})$. Alternatively, one can work with the preimage $\varphi^{-1}(\{0\})$, which yields the *conjunctive normal form* $\varphi(p_1,p_2,p_3) = (p_1 \vee p_2 \vee p_3) \wedge (\neg p_1 \vee p_2 \vee p_3)$ for φ. (How?)

In the preceding case, neither the disjunctive nor the conjunctive normal form yields the simplest expression for φ, namely, $\varphi(p_1,p_2,p_3) = p_2 \vee p_3$, the latter being a *minimal realization* of φ. Readers who have had a course in elementary discrete mathematics will recall that boolean functions play a role in describing how the outputs of certain electrical circuits (with 1 = current flows, and 0 = current does not flow) depend on their inputs (here, boolean functions are often referred to as *switching functions*). In this case, the operations $\wedge$, $\vee$, and $\neg$ correspond to certain *logic gates* (the AND, OR, and NOT gates). When two wires enter an AND (respectively, OR) gate, current flows as an output just when both incoming wires carry a current (respectively, at least one wire carries a current). When a single wire enters a NOT gate, current flows as an output if and only if the entering wire does not carry a current. See Stone (1973) for further details.

In view of the above remarks, it is obviously desirable to have a minimal realization of each boolean function, in order to create circuits that function in a prescribed way using as few gates as possible. Various algorithms have been devised for creating such minimal realizations from an exhaustive description of the boolean function that specifies the desired behavior of the circuit. See, for example, Quine (1955) and McCluskey (1956). Since there are $2^{(2^n)}$ n-ary boolean functions, this task might appear at first glance to involve the cataloguing of an enormous number of minimal realizations. But consider the function φ given above, which has the minimal realization $\varphi(p_1,p_2,p_3) = p_2 \vee p_3$. Suppose that $\sigma : [3] \to [3]$ is a permutation and $\varphi'(p_1,p_2,p_3) := \varphi(p_{\sigma(1)},p_{\sigma(2)},p_{\sigma(3)})$. Then a minimal realization of φ' is given by $\varphi'(p_1,p_2,p_3) = p_{\sigma(2)} \vee p_{\sigma(3)}$.

Suppose that we define an equivalence relation $\sim$ on the set of all functions $\varphi : \{0,1\}^n \to \{0,1\}$ by $\varphi \sim \varphi' \Leftrightarrow$ there exists a permutation $\sigma : [n] \to [n]$ such that $\varphi'(p_1,p_2,\ldots,p_n) = \varphi(p_{\sigma(1)},p_{\sigma(2)},\ldots,p_{\sigma(n)})$. Then, by the preceding observation, we need only find a minimal realization of one representative from each $\sim$ equivalence class. The number of such equivalence classes follows from Pólya's first theorem, where $D = \{0,1\}^n$, $C = \{0,1\}$, and G denotes the permutation group on $\{0,1\}^n$ given by $G = \{g_\sigma : \sigma \text{ is a permutation of } [n] \text{ and } g_\sigma(p_1,p_2,\ldots,p_n) = (p_{\sigma(1)},p_{\sigma(2)},\ldots,p_{\sigma(n)})\}$. Then $\varphi' \sim \varphi$ if and only if $\varphi' = \varphi \circ g$ for some $g \in G$, and so the number of $\sim$ equivalence classes is given by $Z(G;z_1,z_2,\ldots,z_{2^n})|_{z_i \equiv 2}$. You are asked to evaluate this expression when $n = 3$ in Exercise 10.13. As an illustration, a derivation of the number of $\sim$

equivalence classes of 2-ary boolean functions follows. There are $2^{(2^2)} = 16$ boolean functions of two variables. Here G contains two permutations of $\{0,1\}^2$:

(i) $(p_1, p_2) \mapsto (p_1, p_2)$, which has the cycle decomposition (00)(01)(10)(11), and

(ii) $(p_1, p_2) \mapsto (p_2, p_1)$, which has the cycle decomposition (00)(11)(01,10).

So $Z(G; z_1, z_2, z_3, z_4)|_{z_i \equiv 2} = \frac{1}{2}(z_1^4 + z_1^2 z_2)|_{z_1 = z_2 = 2} = 12$. It is a worthwhile exercise to identify the members of each of these 12 equivalence classes.

We can actually cut down on the number of equivalence classes by considering a coarser equivalence relation $\sim^*$, defined as follows. We consider immediately the case of n-ary boolean functions $\varphi : \{0,1\}^n \to \{0,1\}$. Here, $G^* = \{g_{\sigma;s_1,s_2,\ldots,s_n} : \sigma \text{ is a permutation of } [n]$; each of the symbols s_i is either a blank or the negation symbol $\neg$; and $g_{\sigma;s_1,s_2,\ldots,s_n}(p_1, p_2, \ldots, p_n) := (s_1 p_{\sigma(1)}, s_2 p_{\sigma(2)}, \ldots, s_n p_{\sigma(n)})\}$. Note that in this case $|G^*| = n!\, 2^n$. The relation $\varphi' \sim^* \varphi$ holds if and only if there exists a $g \in G^*$ such that $\varphi' = \varphi \circ g$. In Exercise 10.14 you are asked to determine the number of $\sim^*$ equivalence classes of 3-ary boolean functions. As a guide to completing this exercise, here is a derivation of the solution to this problem for 2-ary functions. The eight members of G^* are the following.

(1) $(p_1, p_2) \mapsto (p_1, p_2)$, with cycle decomposition (00)(01)(10)(11);

(2) $(p_1, p_2) \mapsto (\neg p_1, p_2)$, with cycle decomposition (00,10)(01,11);

(3) $(p_1, p_2) \mapsto (p_1, \neg p_2)$, with cycle decomposition (00,01)(10,11);

(4) $(p_1, p_2) \mapsto (\neg p_1, \neg p_2)$, with cycle decomposition (00,11)(01,10);

(5) $(p_1, p_2) \mapsto (p_2, p_1)$, with cycle decomposition (00)(11)(01,10);

(6) $(p_1, p_2) \mapsto (\neg p_2, p_1)$, with cycle decomposition (00,10,11,01);

(7) $(p_1, p_2) \mapsto (p_2, \neg p_1)$, with cycle decomposition (00,01,11,10);

(8) $(p_1, p_2) \mapsto (\neg p_2, \neg p_1)$, with cycle decomposition (00,11)(01)(10);

and so $Z(G^*; z_1, z_2, z_3, z_4)|_{z_i \equiv 2} = \frac{1}{8}(z_1^4 + 2z_1^2 z_2 + 3z_2^2 + 2z_4)|_{z_i \equiv 2} = 6$.

Remark 10.7.1. The hand computation of the cycle index $Z(G^*; z_1, z_2, \ldots, z_{2^n})$ for $n > 3$ is out of the question. Indeed, a computer in the early 1950s required an enormous amount of time just to compute this index for $n = 4$. See Harrison (1971) for a list of these cycle indices for $n \leq 6$.

References

[1] C. Berge (1971): *Principles of Combinatorics*, Academic Press. MR0270922

[2] N. De Bruijn (1959): *Generalization of Pólya's fundamental theorem in enumerative combinatorial analysis*, Indagationes Mathematicae **21**, 59–79. MR0105370

[3] F. Harary and E. Palmer (1973): *Graphical Enumeration*, Academic Press. MR0357214

[4] M. Harrison (1971): *Counting theorems and their applications to switching functions*, in Recent Developments in Switching Theory (A. Mukhopadhyay, ed.), Academic Press. MR0278844

[5] E. McCluskey (1956): *Minimization of boolean functions*, Bell Systems Technical Journal **35**, 1417. MR82876

[6] G. Pólya (1937): *Kombinatorische Anzahlbestimmungen für Gruppen, Graphen, und Chemische Verbindungen*, Acta Mathematica **68**, 145–253. MR1577579

[7] W. Quine (1955): *A way to simplify truth functions*, American Mathematical Monthly **59**, 627–631. MR75886

[8] H. Stone (1973): *Discrete Mathematical Structures and their Applications*, Science Research Associates.

[9] S. Warner (1965): *Modern Algebra, Volume I*, Prentice-Hall. MR1068318

Exercises

10.1. (a) Prove that if G is a permutation group on $[n]$, $\varphi_1, \varphi_2 \in [k]^{[n]}$, and $\varphi_1 \sim \varphi_2$, then, for all $j \in [k]$, $|\varphi_1^{-1}(\{j\})| = |\varphi_2^{-1}(\{j\})|$.
(b) Prove that if $G = S_n$, $\varphi_1, \varphi_2 \in [k]^{[n]}$, and $|\varphi_1^{-1}(\{j\})| = |\varphi_2^{-1}(\{j\})|$ for all $j \in [k]$, then $\varphi_1 \sim \varphi_2$.

10.2. If $n, k \in \mathbb{P}$, simplify $\sum k^{\lambda_1+\lambda_2+\cdots+\lambda_n}/1^{\lambda_1}2^{\lambda_2}\cdots n^{\lambda_n}\lambda_1!\,\lambda_2!\cdots\lambda_n!$, where the sum is taken over all $\lambda_1, \ldots, \lambda_n$ such that $\lambda_1 + 2\lambda_2 + \cdots n\lambda_n = n$ and $\lambda_i \in \mathbb{N}$.
How does this compare to Exercise 2.9?

10.3. (a) Find $Z(D_6; z_1, \ldots, z_6)$, where D_6 is the dihedral group of a regular hexagon.
(b) If the sides of this hexagon may be painted with any of the colors black, gold, red, and green, how many $\sim$ equivalence classes of colorings are there relative to D_6?

10.4. (a) Determine the cycle index $Z(D_3; z_1, z_2, z_3)$ of the dihedral group on an equilateral triangle. If the sides of the triangle can be colored with any of the colors red, white, and blue, how many $\sim$ equivalence classes of colorings are there relative to D_3?
(b) Determine the pattern inventory $P(D_3; r, w, b)$.
(c) Determine $P(D_3; 1, 1, b)$.
(d) Compare $P(D_3; 0, w, b)$ and $P(D_3; 1, w, b)$.
(e) Answer the two questions posed at the end of section 10.3.

10.5. Prove Theorem 10.2.5.

10.6. (a) How many different necklaces with nine beads can be constructed if beads of three different colors are available?
(b) Answer part (a) if each color appears exactly three times.

10.7. Complete the proof of Lemma 10.4.2 by showing that the pairs in (10.4.10) are distinct.

10.8. Complete the proof of Lemma 10.4.3.

10.9. Derive the cycle index of $S_4^{(2)}$ from the cycle index of S_4.

10.10. (a) Let $\{R, S\}$ be a partition of $[n]$ with $|R| = r$ and $|S| = s$. Let H and K be permutation groups, respectively, on R and S, with cycle indices $Z(H; x_1, \ldots, x_r)$ and $Z(K; x_1, \ldots, x_s)$. If $G := \{h \cup k : h \in H \text{ and } k \in K\}$, with functions construed extensionally, G is obviously a permutation group on $[n]$. Determine $Z(G; x_1, \ldots, x_{r+s})$.

(b) Use the results of part (a), along with Pólya theory, to determine the number of ways to distribute four indistinguishable spheres, four indistinguishable cubes, and four indistinguishable cylinders among three distinguishable urns.

(c) Answer part (b) above without using Pólya theory, based on elementary considerations.

10.11. Prove Theorem 10.5.2.

10.12. If Φ is the set of injective colorings $\varphi : [n] \to [k]$ and G is a permutation group on $[n]$, into how many G-classes is Φ partitioned? How many classes are there if $G = S_n$?

10.13. Determine the number of $\sim$ equivalence classes of 3-ary boolean functions.

10.14. Determine the number of $\sim^*$ equivalence classes of 3-ary boolean functions.

10.15. Determine the number of $\approx$ equivalence classes of 3-ary boolean functions, where $\phi_1 \approx \phi_2$ if and only if there exists a permutation σ on [3], a sequence (s_1, s_2, s_3), with each s_i a blank or a logical negation symbol, and a symbol s which is either a blank or a logical negation symbol, such that $\forall (p_1, p_2, p_3) \in \{0,1\}^3$, $\phi_2(p_1, p_2, p_3) = s(\phi_1(s_1 p_{\sigma(1)}, s_2 p_{\sigma(2)}, s_3 p_{\sigma(3)})$. *Hint*: Apply Theorem 10.6.3.

10.16. If G is a permutation group on X and H is a permutation group on Y, we say that G and H are *permutationally equivalent* if there exists a group isomorphism $f : G \to H$ and a bijection $b : X \to Y$ such that, for all $x \in X$ and all $g \in G$, it is the case that $b(g(x)) = f(g)(b(x))$. Prove that if G and H are permutationally equivalent, then their cycle indices are identical.

10.17. Let G be *any* permutation group on $\{0,1\}^{[n]}$. Recall that boolean functions $\varphi, \varphi' : \{0,1\}^n \to \{0,1\}$ are said to be *G-equivalent* (symbolized $\varphi \sim \varphi'$) if there exists a $g \in G$ such that $\varphi' = \varphi \circ g$. Prove that the number of G-equivalence classes must be strictly greater than 2^n.

10.18. Instantiate formula (10.6.4) for the case $|D| = n$ and $G = \{i_D\}$. How is the result related to the solution to Exercise 6.2(b)?