Quick Check

1. Let $n = 2k+1$. Prove that the Eulerian number $A(n, k+1)$ is even.
2. Give a simple proof of the fact that $A(2n, n+1) \geq n!^2$.
3. Find a direct proof for the formula $A(n, 2) = 2^n - n - 1$ without using the general formula for $A(n, k)$.

4.2 The cycle structure of permutations

So far, we have viewed permutations as linear orders of objects. This is not the only way to view them, which makes them even more interesting.

4.2.1 Stirling numbers of the first kind

Let us now turn to the problem of the second tourist bus. Recall that the n tourists traveling on that bus chose to go to a full-service restaurant, where they sat down around k identical tables. Two seating arrangements are considered identical if each person has the same left neighbor in both of them. In how many different ways can the tourists do this? (Note that we have considered this problem before, in Example 3.31, where we solved it using generating functions.)

At first sight, it is not even clear that this problem is related to permutations. Indeed, permutations are, based on what we have learned so far, linear orders. Here, the tables are circular. Furthermore, in permutations, the order of the objects certainly matters, whereas here the tables are all identical. However, we will show that this is very much a permutation enumeration problem; we just need to broaden our conception of permutations to understand that.

Let $p = p_1 p_2 \cdots p_n$ be a permutation. Then we can view p as a *function*, in fact, a bijection, from $[n]$ onto $[n]$, namely, the function defined by $p(i) = p_i$ for all $i \in [n]$. This way of looking at permutations has a few advantages on its own. One of them is that it provides for a natural way of defining the *product* of two permutations.

Definition 4.13 *Let f and g be two permutations, that is, two bijections from $[n]$ onto $[n]$. Then fg is the permutation defined by $fg(i) = g(f(i))$ for all $i \in [n]$.*

In other words, fg is the composition of f and g as functions.

Example 4.14 *Let $n = 6$, and let $f = 421563$, and let $g = 361524$. Then the reader is invited to verify that $fg = 563241$, while $gf = 134625$.*

You may wonder why we set $fg(i) = g(f(i))$ as opposed to $fg(i) = f(g(i))$. The reason for this is that we want to apply the functions in the order in which they are multiplied together. In $g(f(i))$, we *first* apply f to i, and *then* we apply g to $f(i)$.

We would like to stress that, as the above example shows, fg and gf are in general not identical; in other words, multiplication of permutations is not a commutative operation.

Another useful property of permutations as functions is that we can now talk about the *inverse* of a permutation. The inverse of the permutation p, denoted by p^{-1}, is just the inverse of the bijection p. That is, if $p(i) = j$, then $p^{-1}(j) = i$. Equivalently, we have $pp^{-1} = p^{-1}p = 123\cdots n$, the increasing permutation, which we also call the *identity permutation.*

The fact that we are able to multiply permutations together and take their inverses, along with some other important properties, opens new avenues of studying permutations. Many of these aspects belong to *group theory*, where the set of all n-permutations is called the *symmetric group* of degree n and is denoted S_n. Though the theory of permutation groups is beyond the scope of this book, we will sometimes use the notation S_n for the set of all n-permutations.

As we are now able to multiply n-permutations together, we can in particular talk about *powers* of the permutation p. The computation of powers of p is even easier than that of arbitrary products. Indeed, we only need to draw the *short diagram* of p and watch what happens to each element of $[n]$ if p is repeatedly applied. (The short diagram of a function $f : [n] \to [n]$ consists of n dots, representing the elements of $[n]$, and n arrows, representing the action of f on $[n]$.)

Example 4.15 *Let* $p = 321645$*, and let us take a closer look at permutations* p^k*, with* $k = 1, 2, \cdots$*. Using the short diagram of* p *shown in Figure 4.2, we find the following:*

- $p = 321645$,
- $p^2 = 123564$,
- $p^3 = 321456$,
- $p^4 = 123645$,
- $p^5 = 321564$,
- $p^6 = 123456$,
- $p^7 = 321645 = p$, *and so on.*

We see that the entries 1 and 3 keep being permuted among each other, the entry 2 keeps being fixed, and the entries 4, 5, and 6 keep being cyclically permuted among themselves.

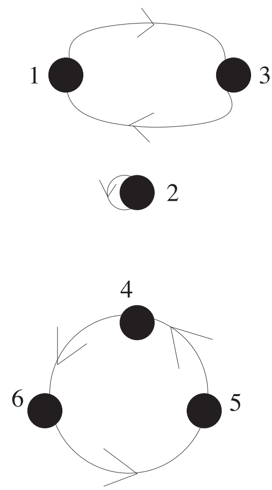

Figure 4.2
The short diagram of $p = 321645$.

In general, if p is an n-permutation, and $i \in [n]$, then there is a positive integer k so that $p^k(i) = i$. Indeed, by the pigeonhole principle, there exist two positive integers t and u so that $t > u$ and $p^t(i) = p^u(i)$ (since there are only n possibilities for $p^t(i)$, but an infinite number of possibilities for t). This implies (why?) that $p^{t-u}(i) = i$, so our claim is proved by setting $k = t - u$.

Let us now choose the *smallest* positive k so that $p^k(i) = i$. Then, if we draw the diagram of p, the entry i will lie on a circle C of length k. As p is a bijection, this circle must be *disjoint* from the rest of the cycles in the short diagram of p. Therefore, if $j \in C$, then $p(j) \in C$, so p permutes the entries of C among each other. Furthermore, $p(j)$ is the left neighbor of j in C, so p permutes the entries of C among each other *cyclically*. Therefore, C is called the *cycle* of p containing i.

If C has fewer than n entries, then let us choose an entry i' that is not part of C. Repeat the above procedure, with i' playing the role of i, to find the cycle C' of p containing i'. Then C' and C are disjoint since p is a bijection. If $C \cup C' = [n]$, then stop, otherwise continue this way until the entire set $[n]$ has been decomposed into disjoint cycles. The obtained set $\{C, C', \cdots\}$ of cycles is called the *cycle decomposition* of p.

The reader may feel that we are cheating as we talk about *the* cycle decomposition while we have not even proved yet that each permutation p has a *unique* cycle decomposition. In order to allay those concerns, we point out that, if we know p, then we know $p(i)$ for all $i \in [n]$, and therefore we can draw the entire diagram of p.

If we want to write a permutation in a way that makes it easy to see what its cycles are, we can put entries in the same cycle within parentheses in the order in which they appear in the cycle.

Example 4.16 *Using cycle notation, the permutation* $p = 321645$ *can be written as* $(31)(2)(645)$.

Hopefully, you are now asking why did I not write (13)(2)(456) or maybe (564)(2)(21). Those notations, and other notations, are also acceptable as long as they all tell correctly what the image of each entry is. (For instance, (456) cannot be replaced by (465), since 4 goes to 5 in the first cycle, while in the second one, 4 goes to 6.) Certainly, the order among the cycles can be changed in any way, and each cycle can be started anywhere. However, for the sake of clarity, it is certainly advantageous to define a *canonical* cycle notation of permutations, which is unique for each permutation.

Definition 4.17 *Let* p *be a permutation. Then the* canonical cycle notation *of* p *consists of writing each cycle of* p *with its largest entry first, then ordering the cycles in increasing order of their largest entries.*

Example 4.18 *The canonical cycle notation of the permutation* p *of Example 4.16 is* $(2)(31)(645)$.

By now the reader might have noticed where we are trying to go with all this. We have shown that a permutation can be conceived not only as a linear order, but also as a set of cycles. This is precisely what we need for the question that started this section, that is, how many ways are there for n tourists to sit down around k identical circular tables. The number that answers this question is so important that it has its own name.

Definition 4.19 *The number of* n*-permutations having exactly* k *cycles is denoted by* $c(n, k)$ *and is called a* signless Stirling number of the first kind.

We promise to explain by the end of the section why the mysterious adjective "signless" is used here. The name "Stirling numbers" may sound a little less surprising. Indeed, the Stirling numbers of the second kind were the numbers of partitions of $[n]$ into k blocks, and now we hear that the (signless) Stirling numbers of the first kind are the numbers of permutations of $[n]$ having k cycles. The connection between these two arrays of numbers is actually a lot stronger than that, as we will explain later in this section as well.

Note that the number of all n-permutations to any number of cycles is, of course, equal to

$$n! = \sum_{k=1}^{n} c(n, k),$$

which demystifies the result of Example 3.31. Recall that, in that example, we found that the number of ways for n people to sit around an unspecified number of circular tables was $n!$. This looked mysterious since at that point in our studies we only knew permutations as linear orders, not as unions of cycles.

By now, the reader is a seasoned veteran of combinatorial definitions and therefore will not be surprised to know that we set $c(n, k) = 0$ for $k > n$, and $c(n, 0) = 0$ for positive n, with $c(0, 0) = 1$.

It is clear that $c(n, n) = 1$, since an n-permutation f can only have n cycles if each of those cycles is of length 1, that is, when $f(i) = i$ for all i. This happens when f is the identity map of $[n]$. The number $c(n, 1)$ is a little bit more complex to determine, but we have essentially done it in Example 1.22. The reader is asked to try to prove that $c(n, 1) = (n - 1)!$ at this point, then compare her argument to ours in the mentioned example.

While an explicit formula for the numbers $c(n, k)$ in the general case exists, it is very complicated, and its proof is beyond the scope of this book. We will see that this is not a very serious problem, since we can find the numbers $c(n, k)$ we want remarkably fast by some other means. To that end, we first prove a triangular recurrence for the numbers $c(n, k)$.

Theorem 4.20 *For all positive integers n and k satisfying $k \leq n$, the recurrence relation*

$$c(n, k) = c(n - 1, k - 1) + (n - 1)c(n - 1, k) \tag{4.6}$$

holds.

We can verify this identity in the special case of $k = 1$. Then we know that $c(n, 1) = (n - 1)!$, while on the right-hand side, $c(n - 1, 0) = 0$ and $c(n - 1, 1) = (n - 2)!$, showing that the identity indeed holds in this special case. Let us now turn to the general case.

Proof: (of Theorem 4.20) All we need to show is that the right-hand side of (4.6) also enumerates permutations of $[n]$ having k cycles. Let p be a permutation counted by $c(n, k)$. Then the entry n of p can either form a cycle by itself or can share a cycle with other entries. (Just as the last tourist entering the restaurant may either sit at a new table or at a table where there are some other tourists.) In the first case, the remaining $n - 1$ entries form a permutation with $k - 1$ cycles, which they can do in $c(n - 1, k - 1)$ ways. In the second case, the remaining entries form a permutation q having k cycles, which they can do in $c(n - 1, k)$ ways. The entry n can be inserted into any existing cycle of q, before any element. Note that inserting q into the end of a cycle would not be different from inserting it into the beginning of that same cycle. All the permutations obtained this way will be different (why?), resulting in $(n - 1)c(n - 1, k)$ permutations. Combining the counts in the two cases completes the proof. $\diamondsuit$

In Exercise 35, the reader will be asked to provide a nongenerating function proof for the mysterious result of Example 3.37. Ideas like the one in the above proof can be useful when solving that exercise.

The following theorem provides a very simple way of determining the numbers $c(n,k)$. This is the method that we promised for computing these numbers.

Theorem 4.21 *For all positive integers n, the identity*

$$\sum_{k=1}^{n} c(n,k)x^k = (x+n-1)(x+n-2)\cdots(x+1)x \tag{4.7}$$

holds.

In other words, if we need the numbers $c(n,k)$ for some fixed n, all we need to do is to expand the product $(x+n-1)(x+n-2)\cdots(x+1)x$, which is a breeze for any software package, and then we can read off the numbers $c(n,k)$ as the coefficients of the obtained polynomial.

Example 4.22 *Let $n = 3$. Then $(x+2)(x+1)x = x^3 + 3x^2 + 2x$, yielding $c(3,3) = 1$, $c(3,2) = 3$, and $c = (3,1) = 2$. We know that these numbers are correct because there is one 3-permutation with three cycles (the identity, since all cycles must be 1-cycles), there are three 3-permutations with two cycles (one cycle has to be a 1-cycle, the other a 2-cycle—we choose the 1-cycle, then must put the remaining two entries in a 2-cycle), and there are two 3-permutations with one cycle (we can choose the image of 1, and the rest follows).*

Proof: (of Theorem 4.21) Note that the right-hand side can be written in the short form $(x+n-1)_n$. We are going to prove our claim by induction on n.

For $n = 1$, our equation reduces to $x = x$, so our claim trivially holds. Now assume that the claim holds for $n-1$, that is,

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

In order to obtain $(x+n-1)_n$, that is, the right-hand side of (4.7), let us multiply both sides of this equation by $x+n-1$. We get

$$\sum_{k=1}^{n-1} c(n-1,k)x^{k+1} + (n-1)\sum_{k=1}^{n-1} c(n-1,k)x^k = (x+n-1)_n.$$

Note on the left-hand side that the coefficient of x^k is precisely $c(n-1,k-1) + (n-1)c(n-1,k)$, which, by Theorem 4.20, is equal to $c(n,k)$. (This is true even in the special case of $k = n$.) Therefore, the last equation reduces to

$$\sum_{k=1}^{n} c(n,k)x^k = (x+n-1)_n,$$

and completes the proof of our induction step. ◊

Note that (4.7) stays true if the sum on the left-hand side is taken from $k = 0$ instead of $k = 1$. In fact, that replacement will only change the value of the left-hand side when $n = 0$, since otherwise $c(n, 0) = 0$.

The time has come for us to reveal what connects the signless Stirling numbers of the first kind to the Stirling numbers of the second kind, and what the adjective "signless" means in the first place.

Recall from your studies in linear algebra that the set V of all polynomials with real coefficients forms a *vector space* over the field of real numbers, and that the infinite set B of polynomials $\{1, x, x^2, \cdots\}$ forms a *basis* of this vector space. In the very unlikely case that you forgot what that means, we remind you that it means that each element of V can be written uniquely as a linear combination of the elements of B, so that the coefficients of that linear combination are real numbers.

What formula (4.7) tells us, in this terminology, is that, when we write the polynomial $(x + n - 1)_n$ as a (unique) linear combination of the elements of B, then the *coefficients* of that linear combination are precisely the signless Stirling numbers of the first kind, in the right order.

Now let us revisit Exercise 11 of Chapter 2, which we are sure the reader has already solved. In that exercise, we proved that

$$x^n = \sum_{k=0}^{n} S(n, k)(x)_k. \tag{4.8}$$

In other words, if we write the elements of B as linear combinations of the elements of the basis $B' = \{1, (x)_1, (x)_2, \cdots\}$ of V, then the coefficients will be precisely the Stirling numbers of the second kind, and in the right order.

Formulae (4.7) and (4.8) do sound quite similar to each other, but one can see that they are not quite "inverses" of each other in the loose sense that one would "undo" what the other does. This is because, while one expresses the polynomials $(x + n - 1)_n$ in terms of the polynomials x^n, the other expresses the polynomials x^n not in terms of the polynomials $(x + n - 1)_n$, but rather in terms of the polynomials $(x)_k$. This will be corrected by the following definition and proposition.

Definition 4.23 *For all nonnegative integers n and k, set $s(n, k) = (-1)^{n-k}c(n, k)$. The numbers $s(n, k)$ are called the* Stirling numbers of the first kind.

This takes care of the adjective "signless" in the name of the numbers $c(n, k)$. This definition also lets us translate (4.7) in the following manner.

Proposition 4.24 *For all nonnegative integers n, we have*

$$\sum_{k=0}^{n} s(n, k)x^k = (x)_n. \tag{4.9}$$