4

Counting permutations

Two buses full of tourists arrived at a large rest area. The n tourists who traveled on the first bus entered the self-service food court one by one (we denote them by elements of $[n]$ in the order in which they get off the bus) and they formed a line at each of the k different concessions that were open. Nobody passed anybody, so the person who got off the bus first got in line first, the person who got off the bus second got in line second, and so on. The n tourists who traveled on the second bus entered a full-service restaurant, where they took their seats around each of k identical circular tables. In how many ways could tourists on each bus proceed?

Both questions asked above are fundamental counting problems of combinatorics. They are both very closely related to permutations, that is, arrangements of the elements of $[n]$ in a line that use each element exactly once. The answers to both questions are also closely related to set partitions, more precisely to the Stirling numbers of the second kind.

4.1 Eulerian numbers

Let us first consider the tourists on the first bus, who formed lines at each of k counters. If we denote the n tourists by the elements of $[n]$ in the order in which they arrived at the food court, then the k lines at the concessions will correspond to k increasing subsequences. (For example, if $n = 6$ and $k = 2$, and the first, fourth, and fifth people go to the first concession, and the second, third, and sixth people go to the second concession, then we get the increasing subsequences 145 and 236.) If we write these k-increasing subsequences one after another in a specified order (say in the order the concessions are located), we get a permutation $p_1p_2\cdots p_n$ of the set $[n]$. In this permutation, there are *at most* $k-1$ entries p_i that are followed by an entry p_{i+1} satisfying $p_i > p_{i+1}$. Indeed, that can only happen if p_{i+1} starts a new increasing subsequence (coming from a line in front of a new concession), and that happens $k-1$ times.

This way of decomposing permutations into increasing subsequences of consecutive elements is our main topic in this section. Therefore, we make the concept formal.

DOI: 10.1201/9781003304272-4

Definition 4.1 *If k is the smallest integer so that the permutation $p = p_1 p_2 \cdots p_n$ can be decomposed into k increasing subsequences of consecutive entries, then we say that p has k* ascending runs*, or, if there is no danger of confusion,* runs*.*

Example 4.2 *The permutation $p = 245169378$ has three ascending runs, namely, 245, 169, and 378.*

So the tourists of the first bus form a permutation with at most k ascending runs.

Sometimes it is easier to work not with the ascending runs themselves, but with the breakpoints between them. Clearly, a new ascending run must start every time a larger element is followed by a smaller one. This is again a basic phenomenon, and it has its own name.

Definition 4.3 *Let $p = p_1 p_2 \cdots p_n$ be a permutation. We say that $i \in [n-1]$ is a* descent *of p if $p_i > p_{i+1}$. If $i \in [n-1]$ is not a descent, then it is called an* ascent*.*

Note that the descents are the *positions* of p, not its entries.

Example 4.4 *The permutation $p = 245169378$ has two descents, $i = 3$ and $i = 6$. Therefore, p has six ascents.*

At this point, the reader should stop and verify that p has $k - 1$ descents if and only if p has k ascending runs. Figure 4.1, showing the permutation of Example 4.4, should help the reader visualize this fact.

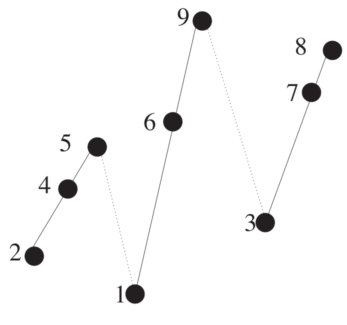

Figure 4.1
Ascending runs and descents in a permutation.

For the rest of this book, permutations of the set $[n]$ will be called n-permutations. If we could tell how many n-permutations there are with exactly k ascending runs, we could also tell how many n-permutations there are with *at most* k ascending runs by the addition principle. It turns out that this is a good way to go, because counting permutations with exactly k ascending runs is easier than counting permutations with at most k ascending runs. Therefore, it is the former to which we give a name.

Definition 4.5 *Let k and n be positive integers satisfying $k \leq n$. Then the number of n-permutations having exactly k ascending runs is denoted by $A(n,k)$ and is called an* Eulerian number.

Example 4.6 *There are four 3-permutations with exactly two ascending runs. They are 132, 213, 231, and 312. Therefore, $A(3,2) = 4$.*

Example 4.7 *For all n, there is one permutation (the increasing one) with one ascending run, and there is one permutation (the decreasing one) with n ascending runs. Therefore, $A(n,1) = A(n,n) = 1$.*

We can extend the definition of $A(n,k)$ for other nonnegative integers by setting $A(n,0) = 0$ and $A(n,k) = 0$ for $k > n$.

An explicit formula for Eulerian numbers has been known since the late nineteenth century. However, we present a relatively recent and remarkably simple proof that was communicated to this author by Richard Stanley in 2003.

Theorem 4.8 *For all nonnegative integers n and k satisfying $k \leq n$, the Eulerian numbers are obtained by the explicit formula*

$$A(n,k) = \sum_{i=0}^{k} (-1)^i \binom{n+1}{i} (k-i)^n. \tag{4.1}$$

Before we start the proof, let us warm up by looking at the special cases of $k = 1$ and $k = 2$. If $k = 1$, then the permutation must have zero descents, so it must be the increasing permutation, implying that $A(n,1) = 1$. The reader can verify that (4.1) holds in this special case. If $k = 2$, then our permutations have two ascending runs. That means that, in order to obtain these permutations, we need to split $[n]$ into nonempty sets X and Y (the future ascending runs) in one of $2^n - 2$ ways, and then order these two sets increasingly, and finally concatenate the obtained strings. We have to be careful, however. If the largest entry of X is smaller than the smallest entry of Y (in other words, when $X = [j]$ and $Y = \{j+1, j+2, \cdots, n\}$), then the obtained permutation will be $12 \cdots n$, which has only one run. This happens in $n-1$ cases, showing that $A(n,k) = 2^n - 2 - (n-1) = 2^n - n - 1$, which is in concordance with (4.1) and also with Example 4.6.

Now we start seeing why proving (4.1) is not easy. As k grows, there could be more and more situations when, in certain circumstances, two substrings that are supposed to form two ascending runs form only one, like X and Y did above. This is a serious problem, because, on one hand, we get fewer ascending runs than we wanted, and, on the other hand, some permutations are obtained more than once. Fortunately, our new friend the inclusion-exclusion principle, as well as an ingenious argument, will help us out.

Proof: (of Theorem 4.8) Let us say that the k concessions are separated by $k-1$ fences, and the tourists just fill the spaces between the fences, as in

$$p = 245|136||79|8. \tag{4.2}$$

Every tourist is free to go to any of the k concessions, which creates k^n possibilities. However, many of these will not be good for our purposes since they will result in a permutation with fewer than k runs. There are two reasons for this that could happen.

(a) There could be empty concessions (no entries between two fences), and

(b) it could happen that the entry immediately preceding a fence is smaller than the following entry.

The reader is invited to verify the easy statements that if neither of (a) and (b) occurs, then we get an n-permutation with exactly k runs, and, crucially, we get each such permutation exactly once.

Now we are going to count the obtained permutations in which at least one of (a) and (b) happens, that is, in which something goes wrong. This may look like a daunting task, since we have to keep track of two properties, not just one. Fortunately, we are able to do that simultaneously, thanks to the definition that follows.

Let p be an n-permutation with $k-1$ fences inserted among its entries (possibly at the beginning or the end). We will call this structure, that is, the permutation and the fences together, an *arrangement*. We say that a fence f is *removable* if it satisfies both of the following requirements:

1. If we remove f, we still have a permutation whose entries increase between any two consecutive fences, and
2. the fence f is not immediately followed by another fence.

Note that, for instance, a fence at the very end of a permutation is always removable. For example, in (4.2), the third fence is removable, and the others are not.

If an arrangement has no removable fences, then the corresponding permutation has k ascending runs, and we are happy. If an arrangement does have at least one removable fence, then the corresponding permutation has less than k ascending runs. Using the inclusion-exclusion principle, we will now enumerate these permutations.

Let us call the $n-1$ spaces between two consecutive entries of our permutations and the space at the beginning and at the end of them *positions*. So there are $n+1$ positions, and we will number them from left to right. Let $S \subseteq [n+1]$, and let A_S be the set of arrangements in which there is a removable fence in each position that belongs to S. Note that while there might be more than one fence in any given position, only one of them (the last one) can be removable.

Fortunately, the size of A_S is easy to determine.

Proposition 4.9 *Let $S \subseteq [n+1]$ so that $|S| = i \leq k-1$. Then*

$$|A_S| = (k-i)^n.$$

Proof: First put down $k-i-1$ fences, and line up the n entries to get an arrangement with at most $k-i$ ascending runs. This can be done in $(k-i)^n$ ways. Now insert the remaining i fences in those positions belonging to S. (If there are already fences there, just put the new fence on the right of the old ones in the same position.) This assures that we get removable fences in the positions that belong to S. Moreover, we will get each arrangement with the required property exactly once. Indeed, if a is an arrangement with the required property, then, if we remove the last fence from each position of a that belongs to S, we get the unique original arrangement that leads to a. $\diamondsuit$

Note that we did not need any specific property of S besides its size, implying that

$$R_i = \sum_{\substack{S \subseteq [n+1] \\ |S|=i}} |A_S| = \binom{n+1}{i}(k-i)^n. \tag{4.3}$$

Now we are ready to use the inclusion-exclusion principle. With a slight abuse of notation, we write A_i for the set of arrangements that have a removable fence in position i. Then Theorem 2.38 yields

$$\begin{aligned} |A_1 \cup A_2 \cup \cdots \cup A_{n+1}| &= R_1 - R_2 + R_3 - \cdots + (-1)^k R_{k-1} \\ &= \sum_{i=1}^{k-1} (-1)^{i+1} \binom{n+1}{i}(k-i)^n. \end{aligned}$$

Note that $R_i = 0$ if $i \geq k$ since there are altogether only $k-1$ fences, so there cannot be more than $k-1$ removable fences.

Recalling that the total number of arrangements we could have is k^n, the subtraction principle shows that the number of arrangements with no removable fences is

$$A(n,k) = \sum_{i=0}^{k-1} (-1)^i \binom{n+1}{i}(k-i)^n, \tag{4.4}$$

which was to be proved. $\diamondsuit$

In many aspects, Eulerian numbers behave similarly to binomial coefficients, and in many other aspects, they are closely related to Stirling numbers of the second kind.

An example for the former is *symmetry*, that is, the property, that for any fixed n and any $k \leq n$, the equality $A(n,k) = A(n, n+1-k)$ holds. The reader should try to prove this identity first, then check our two solutions given in Exercise 1.

Just like binomial coefficients and Stirling numbers of the second kind, Eulerian numbers also satisfy a triangular recurrence.

Theorem 4.10 *Let k and n be nonnegative integers satisfying $k \leq n$. Then*

$$A(n,k) = kA(n-1,k) + (n-k+1)A(n-1,k-1).$$

Proof: The left-hand side clearly counts n-permutations with $k-1$ descents; all we need to show is that the right-hand side counts the same.

Let p be a permutation counted by $A(n,k)$. Let p' be the permutation obtained from p by omitting the entry n. Then there are two cases, namely, that p' has as many descents as p, or it has one less. We will show that the two terms of the right-hand side count permutations that fall into these respective cases.

(Case 1) When the omission of the entry n leaves the number of descents of p unchanged.

This happens when n is at the end of p, or when n is inserted right after the ith entry of p' where i is a descent of p'. (Try $\ldots 3n1 \ldots$.) In either case, we are left with a permutation p' that is of length $n-1$ and has $k-1$ descents. There are, by definition, $A(n-1,k)$ such permutations. Each of them is obtained this way from k different permutations p. Indeed, to get these permutations, we can insert n into any of the $k-1$ descents of p', or we can put it to the end of p'. Therefore, $kA(n-1,k)$ permutations fall into this case.

(Case 2) When the omission of the entry n causes the number of descents of p to decrease by 1.

This happens when n is the first entry of p, or when n is inserted right after the ith entry of p' where i is an ascent of p'. (Try $\ldots 1n3 \ldots$.) In either case, we are left with a permutation p' that is of length $n-1$ and has $k-2$ descents. There are $A(n-1,k-1)$ such permutations. Each of them is obtained this way from $n-k+1$ different permutations p. Indeed, to get these permutations, we can insert n into any of the $n-k$ ascents of p', or we can put it to the front of p'. By the product formula, this explains the second term of the right-hand side.

We have shown that the two sides count the elements of the same set; therefore, they have to be equal. $\Diamond$

Theorem 4.10 enables us to prove that the Eulerian numbers occur in many other contexts. We show one example for that here, and leave some others to the exercises.

Example 4.11 *A city has n different hotels, labeled from 1 to n, where 1 denotes the cheapest hotel and n denotes the most expensive one. On each of n consecutive days, one new, mathematically inclined tourist arrives in the city. These tourists arrive in increasing order of spending power. That is, the tourist who arrives on day i can afford to go to any hotels labeled by an element of $[i]$.*

Let $H(n,k)$ be the number of ways the tourists can make their selections so that at the end there will be exactly k hotels that host at least one of our n tourists. Then $H(n,k) = A(n,k)$.

For instance, if $n = 2$, then either both tourists go to hotel 1, or the first goes to 1 and the second goes to 2, yielding $H(2,1) = H(2,2) = 1$, in concordance with $A(2,1) = A(2,2) = 1$. The reader is invited to verify the result for $n = 3$. Note that it is not surprising that the total number of possibilities is $\sum_{k=1}^{n} A(n,k) = n!$, since tourist k has k possibilities.

Solution: (of Example 4.11) We prove the statement by induction on n, the initial case of $n = 1$ being trivial. Assume that we know the statement for $n-1$, and prove it for n. We claim that

$$H(n,k) = kH(n-1,k) + (n-k+1)H(n-1,k-1). \tag{4.5}$$

The left-hand side is, by definition, the number of all possible selections made by the tourists in which k hotels get at least one customer from this group. We claim that the right-hand side counts the objects. Indeed, there are two different ways that a situation counted by the left-hand side could occur, namely, the last tourist either went to a hotel where there was already someone from the group, or not. In the first case, the first $n-1$ tourists split into k hotels in one of $H(n-1,k)$ ways, then the last tourist chooses one of these k hotels. This explains the first term of the right-hand side. In the second case, the first $n-1$ tourists split into $k-1$ hotels in one of $H(n-1,k-1)$ ways, and then the last tourist chooses one of the remaining $n-k+1$ hotels. This explains the second term of the right-hand side, and therefore proves (4.5).

By our induction hypothesis, we know that $A(n-1,k) = H(n-1,k)$, and also that $A(n-1,k-1) = H(n-1,k-1)$. Looking at Theorem 4.10, we see that this means that $A(n,k)$ and $H(n,k)$ are obtained from identical numbers, by identical operations. Therefore, they must be equal. $\diamondsuit$

Note that we proved Theorem 4.10 by a particular kind of induction, that is, we proved that two arrays of numbers have to be identical *because they start the same way* and *they satisfy the same recursive relation*. We will see further applications of this proof technique soon.

The following concept ties descents to another interesting permutation statistic that will be discussed in the following section. A *desarrangement* is a permutation whose first ascent occurs in an even position. We make a single exception to this rule. If n is even, then we will also consider the *decreasing* permutation $n(n-1)\cdots 21$ a desarrangement, as if it had an ascent at its last position n. So 213, 43215 and 4321 are all desarrangements, but 12, 321, and 52134 are not. What interests us here is the *number* of desarrangements of length n.

The term "desarrangements" was coined by Michelle Wachs, to honor her co-author Jacques Désarmenien in an amusing way. The editors of the journal in which the term was first used originally wanted to change it to "disarrangements," which *is* an English word, but when they got the joke, they left "desarrangements" unchanged.

Lemma 4.12 *The number of desarrangements of length n is*

$$J(n) = n!\sum_{i=0}^{n} \frac{(-1)^i}{i!}.$$

Proof: Let n be fixed, and let b_i be the number of n-permutations whose first ascent occurs in position $2i$. In such a permutation, there are $2i$ different possibilities for the relative order of the first $2i+1$ entries. Indeed, the entry in position $2i+1$ can be anything but the smallest of these $2i+1$ entries, but once that entry is chosen the remaining $2i$ entries must all be in decreasing order. Therefore, for $i < n/2$, the ratio of permutations enumerated by b_i to all n-permutations is $\frac{2i}{(2i+1)!}$; consequently

$$b_i = \frac{2i}{(2i+1)!}n! = \left(\frac{1}{(2i)!} - \frac{1}{(2i+1)!}\right) n!.$$

Taking the sum of all possible b_i, we get

$$J(n) = \sum_{i=1}^{\lfloor (n-1)/2 \rfloor} n!\left(\frac{1}{(2i)!} - \frac{1}{(2i+1)!}\right) = n!\left(\frac{1}{2!} - \frac{1}{3!} + \frac{1}{4!} - \cdots\right).$$

This is equivalent to our claim, since our claim was that

$$J(n) = n!\left(1 - 1 + \frac{1}{2} - \frac{1}{6} + \cdots\right).$$

What we proved is the same, without the first two terms that cancel each other anyway. ◊

Finally, we mention that, instead of simply counting permutations with a given number of descents, we could count permutations whose *set* of descents is a given set S, or whose set of descents is contained in a given set S. The reader is invited to look at Exercises 2–6 for results of this kind.