

$$\begin{aligned}
\Pr\{E_2 \cap E_1\} &= \Pr\{E_2 \mid E_1\} \Pr\{E_1\} \\
&= \frac{1}{n-i+1} \cdot \frac{(n-i+1)!}{n!} \\
&= \frac{(n-i)!}{n!}.
\end{aligned}$$

**Termination:** At termination,  $i = n + 1$ , and we have that the subarray  $A[1..n]$  is a given  $n$ -permutation with probability  $(n - (n + 1) + 1)/n! = 0!/n! = 1/n!$ .

Thus, RANDOMIZE-IN-PLACE produces a uniform random permutation. ■

A randomized algorithm is often the simplest and most efficient way to solve a problem. We shall use randomized algorithms occasionally throughout this book.

## Exercises

### 5.3-1

Professor Marceau objects to the loop invariant used in the proof of Lemma 5.5. He questions whether it is true prior to the first iteration. He reasons that we could just as easily declare that an empty subarray contains no 0-permutations. Therefore, the probability that an empty subarray contains a 0-permutation should be 0, thus invalidating the loop invariant prior to the first iteration. Rewrite the procedure RANDOMIZE-IN-PLACE so that its associated loop invariant applies to a nonempty subarray prior to the first iteration, and modify the proof of Lemma 5.5 for your procedure.

### 5.3-2

Professor Kelp decides to write a procedure that produces at random any permutation besides the identity permutation. He proposes the following procedure:

PERMUTE-WITHOUT-IDENTITY( $A$ )

```

1   $n = A.length$ 
2  for  $i = 1$  to  $n - 1$ 
3      swap  $A[i]$  with  $A[\text{RANDOM}(i + 1, n)]$ 

```

Does this code do what Professor Kelp intends?

### 5.3-3

Suppose that instead of swapping element  $A[i]$  with a random element from the subarray  $A[i..n]$ , we swapped it with a random element from anywhere in the array:

PERMUTE-WITH-ALL( $A$ )

```

1   $n = A.length$ 
2  for  $i = 1$  to  $n$ 
3      swap  $A[i]$  with  $A[\text{RANDOM}(1, n)]$ 

```

Does this code produce a uniform random permutation? Why or why not?

#### 5.3-4

Professor Armstrong suggests the following procedure for generating a uniform random permutation:

PERMUTE-BY-CYCLIC( $A$ )

```

1   $n = A.length$ 
2  let  $B[1..n]$  be a new array
3   $offset = \text{RANDOM}(1, n)$ 
4  for  $i = 1$  to  $n$ 
5       $dest = i + offset$ 
6      if  $dest > n$ 
7           $dest = dest - n$ 
8       $B[dest] = A[i]$ 
9  return  $B$ 

```

Show that each element  $A[i]$  has a  $1/n$  probability of winding up in any particular position in  $B$ . Then show that Professor Armstrong is mistaken by showing that the resulting permutation is not uniformly random.

#### 5.3-5 ★

Prove that in the array  $P$  in procedure PERMUTE-BY-SORTING, the probability that all elements are unique is at least  $1 - 1/n$ .

#### 5.3-6

Explain how to implement the algorithm PERMUTE-BY-SORTING to handle the case in which two or more priorities are identical. That is, your algorithm should produce a uniform random permutation, even if two or more priorities are identical.

#### 5.3-7

Suppose we want to create a *random sample* of the set  $\{1, 2, 3, \dots, n\}$ , that is, an  $m$ -element subset  $S$ , where  $0 \leq m \leq n$ , such that each  $m$ -subset is equally likely to be created. One way would be to set  $A[i] = i$  for  $i = 1, 2, 3, \dots, n$ , call RANDOMIZE-IN-PLACE( $A$ ), and then take just the first  $m$  array elements. This method would make  $n$  calls to the RANDOM procedure. If  $n$  is much larger than  $m$ , we can create a random sample with fewer calls to RANDOM. Show that

words of lengths  $l_1, l_2, \dots, l_n$ , measured in characters. We want to print this paragraph neatly on a number of lines that hold a maximum of  $M$  characters each. Our criterion of “neatness” is as follows. If a given line contains words  $i$  through  $j$ , where  $i \leq j$ , and we leave exactly one space between words, the number of extra space characters at the end of the line is  $M - j + i - \sum_{k=i}^j l_k$ , which must be nonnegative so that the words fit on the line. We wish to minimize the sum, over all lines except the last, of the cubes of the numbers of extra space characters at the ends of lines. Give a dynamic-programming algorithm to print a paragraph of  $n$  words neatly on a printer. Analyze the running time and space requirements of your algorithm.

### 15-5 Edit distance

In order to transform one source string of text  $x[1..m]$  to a target string  $y[1..n]$ , we can perform various transformation operations. Our goal is, given  $x$  and  $y$ , to produce a series of transformations that change  $x$  to  $y$ . We use an array  $z$ —assumed to be large enough to hold all the characters it will need—to hold the intermediate results. Initially,  $z$  is empty, and at termination, we should have  $z[j] = y[j]$  for  $j = 1, 2, \dots, n$ . We maintain current indices  $i$  into  $x$  and  $j$  into  $z$ , and the operations are allowed to alter  $z$  and these indices. Initially,  $i = j = 1$ . We are required to examine every character in  $x$  during the transformation, which means that at the end of the sequence of transformation operations, we must have  $i = m + 1$ .

We may choose from among six transformation operations:

- Copy** a character from  $x$  to  $z$  by setting  $z[j] = x[i]$  and then incrementing both  $i$  and  $j$ . This operation examines  $x[i]$ .
- Replace** a character from  $x$  by another character  $c$ , by setting  $z[j] = c$ , and then incrementing both  $i$  and  $j$ . This operation examines  $x[i]$ .
- Delete** a character from  $x$  by incrementing  $i$  but leaving  $j$  alone. This operation examines  $x[i]$ .
- Insert** the character  $c$  into  $z$  by setting  $z[j] = c$  and then incrementing  $j$ , but leaving  $i$  alone. This operation examines no characters of  $x$ .
- Twiddle** (i.e., exchange) the next two characters by copying them from  $x$  to  $z$  but in the opposite order; we do so by setting  $z[j] = x[i + 1]$  and  $z[j + 1] = x[i]$  and then setting  $i = i + 2$  and  $j = j + 2$ . This operation examines  $x[i]$  and  $x[i + 1]$ .
- Kill** the remainder of  $x$  by setting  $i = m + 1$ . This operation examines all characters in  $x$  that have not yet been examined. This operation, if performed, must be the final operation.

As an example, one way to transform the source string `algorithm` to the target string `altruistic` is to use the following sequence of operations, where the underlined characters are  $x[i]$  and  $z[j]$  after the operation:

Operation	$x$	$z$
<i>initial strings</i>	<u>a</u> lgorithm	_
copy	a <u>l</u> gorithm	a_
copy	al <u>g</u> orithm	al_
replace by t	alg <u>o</u> rithm	alt_
delete	algor <u>u</u> ithm	alt_
copy	algori <u>t</u> h	altr_
insert u	algori <u>u</u> th	altru_
insert i	algori <u>i</u> th	altrui_
insert s	algori <u>s</u> th	altruis_
twiddle	algori <u>h</u> th	altruisti_
insert c	algori <u>h</u> th	altruistic_
kill	algori <u>h</u> th_	altruistic_

Note that there are several other sequences of transformation operations that transform `algorithm` to `altruistic`.

Each of the transformation operations has an associated cost. The cost of an operation depends on the specific application, but we assume that each operation's cost is a constant that is known to us. We also assume that the individual costs of the copy and replace operations are less than the combined costs of the delete and insert operations; otherwise, the copy and replace operations would not be used. The cost of a given sequence of transformation operations is the sum of the costs of the individual operations in the sequence. For the sequence above, the cost of transforming `algorithm` to `altruistic` is

$$(3 \cdot \text{cost}(\text{copy})) + \text{cost}(\text{replace}) + \text{cost}(\text{delete}) + (4 \cdot \text{cost}(\text{insert})) \\ + \text{cost}(\text{twiddle}) + \text{cost}(\text{kill}) .$$

- a.** Given two sequences  $x[1..m]$  and  $y[1..n]$  and set of transformation-operation costs, the *edit distance* from  $x$  to  $y$  is the cost of the least expensive operation sequence that transforms  $x$  to  $y$ . Describe a dynamic-programming algorithm that finds the edit distance from  $x[1..m]$  to  $y[1..n]$  and prints an optimal operation sequence. Analyze the running time and space requirements of your algorithm.

The edit-distance problem generalizes the problem of aligning two DNA sequences (see, for example, Setubal and Meidanis [310, Section 3.2]). There are several methods for measuring the similarity of two DNA sequences by aligning them. One such method to align two sequences  $x$  and  $y$  consists of inserting spaces at