# Permutation Generation Methods* 

## ROBERT SEDGEWICK

Program in Computer Science and Division of Appleed Mathematics Brown University, Providence, Rhode Island 02912


#### Abstract

This paper surveys the numerous methods that have been proposed for permutation enumeration by computer. The various algorithms which have been developed over the years are described in detail, and implemented in a modern Algol-like language. All of the algorithms are derived from one simple control structure. The problems involved with implementing the best of the algorithms on real computers are treated in detail. Assembly-language programs are derived and analyzed fully. The paper is intended not only as a survey of permutation generation methods, but also as a tutorial on how to compare a number of different algorithms for the same task


Key Words and Phrases: permutations, combinatorial algorithms, code optimızation, analysis of algorthms, lexicographic ordering, random permutations, recursion, cyclic rotation.

CR Categorles: 3.15, 4.6, 525, 5.30.

## INTRODUCTION

Over thirty algorithms have been published during the past twenty years for generating by computer all $N$ ! permutations of $N$ elements. This problem is a nontrivial example of the use of computers in combinatorial mathematics, and it is interesting to study because a number of different approaches can be compared. Surveys of the field have been published previously in 1960 by D. H. Lehmer [26] and in 1970-71 by R. J. Ord-Smith [29, 30]. A new look at the problem is appropriate at this time because several new algorithms have been proposed in the intervening years.

Permutation generation has a long and distinguished history. It was actually one of the first nontrivial nonnumeric problems to be attacked by computer. In 1956, C. Tompkins wrote a paper [44] describing a number of practical areas where permu-

[^0]tation generation was being used to solve problems. Most of the problems that he described are now handled with more sophisticated techniques, but the paper stimulated interest in permutation generation by computer per se. The problem is simply stated, but not easily solved, and is often used as an example in programming and correctness. (See, for example, [6]).

The study of the various methods that have been proposed for permutation generation is still very instructive today because together they illustrate nicely the relationship between counting, recursion, and iteration. These are fundamental concepts in computer science, and it is useful to have a rather simple example which illustrates so well the relationships between them. We shall see that algorithms which seem to differ markedly have essentially the same structure when expressed in a modern language and subjected to simple program transformations. Many readers may find it surprising to discover that "top-down" (recursive) and "bottom-up"

[^1]
## CONTENTS

INTRODUCTION
1 METHODS BASED ON EXCHANGES
Recursive methods
Adjacent exchanges
Factorial counting
"Loopless" algonthms
Another iterative method
2 OTHER TYPES OF ALGORITHMS
Nested cyclung
Lexicographic algorithms
Random permutations
3 IMPLEMENTATION AND ANALYSIS
A recursive method (Heap)
An terative method (Ives)
A cyclic method (Langdon)
CONCLUSION
ACKNOWLEDGMENTS
REFERENCES
(iterative) design approaches can lead to the same program.

Permutation generation methods not only illustrate programming issues in high-level (procedural) languages; they also illustrate implementation issues in low-level (assembly) languages. In this paper, we shall try to find the fastest possible way to generate permutations by computer. To do so, we will need to consider some program "optimization" methods (to get good implementations) and some mathematical analyses (to determine which implementation is best). It turns out that on most computers we can generate each permutation at only slightly more than the cost of two store instructions.

In dealing with such a problem, we must be aware of the inherent limitations. Without computers, few individuals had the patience to record all 5040 permutations of 7 elements, let alone all 40320 permutations of 8 elements, or all 362880 permutations of 9 elements. Computers
help, but not as much as one might think. Table 1 shows the values of $N$ ! for $N \leq 17$ along with the time that would be taken by a permutation generation program that produces a new permutation each microsecond. For $N>25$, the time required is far greater than the age of the earth!

For many practical applications, the sheer magnitude of $N$ ! has led to the development of "combinatorial search" procedures which are far more efficient than permutation enumeration. Techniques such as mathematical programming and backtracking are used regularly to solve optimization problems in industrial situations, and have led to the resolution of several hard problems in combinatorial mathematics (notably the four-color problem). Full treatment of these methods would be beyond the scope of this paper they are mentioned here to emphasize that, in practice, there are usually alternatives to the "brute-force" method of generating permutations. We will see one example of how permutation generation can sometimes be greatly improved with a backtracking technique.

In the few applications that remain where permutation generation is really required, it usually doesn't matter much which generation method is used, since the cost of processing the permutations far

TABLE 1. Approximate Time Needed to Generate all Permutations of $N$ ( $1 \mu \mathrm{sec}$ per permutation)

| $N$ | $N!$ | Time |
| ---: | ---: | :---: |
|  |  |  |
| 1 | 1 |  |
| 2 | 2 |  |
| 3 | 6 |  |
| 4 | 24 |  |
| 5 | 120 |  |
| 6 | 720 |  |
| 7 | 5040 |  |
| 8 | 40320 |  |
| 9 | 362880 |  |
| 10 | 3628800 | 3 seconds |
| 11 | 39916800 | 40 seconds |
| 12 | 479001600 | 8 minutes |
| 13 | 6227020800 | 2 hours |
| 14 | 87178291200 | 1 day |
| 15 | 1307674368000 | 2 weeks |
| 16 | 20922789888000 | 8 months |
| 17 | 355689428096000 | 10 years |

exceeds the cost of generating them. For example, to evaluate the performance of an operating system, we might want to try all different permutations of a fixed set of tasks for processing, but most of our time would be spent simulating the processing, not generating the permutations. The same is usually true in the study of combinatorial properties of permutations, or in the analysis of sorting methods. In such applications, it can sometimes be worthwhile to generate "random" permutations to get results for a typical case. We shall examine a few methods for doing so in this paper.

In short, the fastest possible permutation method is of limited importance in practice. There is nearly always a better way to proceed, and if there is not, the problem becomes really hopeless when $N$ is increased only a little.

Nevertheless, permutation generation provides a very instructive exercise in the implementation and analysis of algorithms. The problem has received a great deal of attention in the literature, and the techniques that we learn in the process of carefully comparing these interesting algorithms can later be applied to the perhaps more mundane problems that we face from day to day.

We shall begin with simple algorithms that generate permutations of an array by successively exchanging elements; these algorithms all have a common control structure described in Section 1. We then will study a few older algorithms, including some based on elementary operations other than exchanges, in the framework of this same control structure (Section 2). Finally, we shall treat the issues involved in the implementation, analysis, and "optimization" of the best of the algorithms (Section 3).

## 1. METHODS BASED ON EXChANGES

A natural way to permute an array of elements on a computer is to exchange two of its elements. The fastest permutation algorithms operate in this way: All $N$ ! permutations of $N$ elements are produced by a sequence of $N!-1$ exchanges. We shall use the notation

$$
\mathrm{P}[1]:=: \mathrm{P}[2]
$$

to mean "exchange the contents of array elements P[1] and P[2]". This instruction gives both arrangements of the elements $\mathrm{P}[1], \mathrm{P}[2]$ (i.e., the arrangement before the exchange and the one after). For $N=3$, several different sequences of five exchanges can be used to generate all six permutations, for example

$$
\begin{aligned}
& \mathrm{P}[1]=: \mathrm{P}[2] \\
& \mathrm{P}[2]:=: \mathrm{P}[3] \\
& \mathrm{P}[1]=: \mathrm{P}[2] \\
& \mathrm{P}[2]=: \mathrm{P}[3] \\
& \mathrm{P}[1]:=\cdot \mathrm{P}[2] .
\end{aligned}
$$

If the initial contents of P [1] $\mathrm{P}[2] \mathrm{P}[3]$ are A B C, then these five exchanges will produce the permutations $\mathrm{BA} \mathrm{C}, \mathrm{B} \mathrm{CA}$, C B A, C A B, and A C B.
It will be convenient to work with a more compact representation describing these exchange sequences. We can think of the elements as passing through "permutation networks" which produce all the permutations. The networks are comprised of "exchange modules" such as that shown in Diagram 1 which is itself the


## Diagram 1

permutation network for $N=2$. The network of Diagram 2 implements the exchange sequence given above for $N=3$. The elements pass from right to left, and a new permutation is available after each exchange. Of course, we must be sure that the internal permutations generated are distinct. For $N=3$ there are $3^{5}=243$ possible networks with five exchange modules, but only the twelve shown in Fig. 1 are "legal" (produce sequences of distinct permutations). We shall most often represent networks as in Fig. 1, namely drawn vertically, with elements passing from top to bottom, and with the permutation se-



Figure 1. Legal permutation networks for three elements.
quences that are generated explicitly written out on the right.
It is easy to see that for larger $N$ there will be large numbers of legal networks. The methods that we shall now examine will show how to systematically construct networks for arbitrary $N$. Of course, we are most interested in networks with a sufficiently simple structure that their exchange sequences can be conveniently implemented on a computer.

## Recursive Methods

We begin by studying a class of permutation generation methods that are very simple when expressed as recursive programs. To generate all permutations of $\mathrm{P}[1], \cdots, \mathrm{P}[N]$, we repeat $N$ times the step: "first generate all permutations of $\mathrm{P}[1], \cdots, \mathrm{P}[N-1]$, then exchange $\mathrm{P}[N]$ with one of the elements $\mathrm{P}[1], \cdots, \mathrm{P}[N-1]$ ". As this is repeated, a new value is put into $P[N]$ each time. The various methods differ in their approaches to filling $\mathrm{P}[N]$ with the $N$ original elements.
The first and seventh networks in Fig. 1 operate according to this discipline. Recur-


Diagram 3.
sively, we can build up networks for four elements from one of these. For example, using four copies of the first network in Fig. 1, we can build a network for $N=4$, as shown in Diagram 3. This network fills $P[4]$ with the values $D, C, B, A$ in decreasing alphabetic order (and we could clearly build many similar networks which fill $\mathrm{P}[4]$ with the values in other orders).

The corresponding network for five elements, shown in Diagram 4, is more complicated. (The empty boxes denote the network of Diagram 3 for four elements). To get the desired decreasing sequence in $P[5]$, we must exchange it successively with $\mathrm{P}[3], \mathrm{P}[1], \mathrm{P}[3], \mathrm{P}[1]$ in-between generating all permutations of $\mathrm{P}[1], \cdots, \mathrm{P}[4]$.

In general, we can generate all permutations of $N$ elements with the following recursive procedure:

```
Algorithm 1.
    procedure permutations(N);
        begin \(\mathrm{c}=1\);
            toop:
                                    if \(N>2\) then permutations \((N-1)\)
                        endif;
            while \(c<N\) :
                        \(\mathrm{P}[\mathrm{B}[\mathrm{N}, \mathrm{c}]]:=: \mathrm{P}[N] ;\)
                        \(c:=c+1\)
            repeat
    end;
```

This program uses the looping control construct loop $\cdots$ while $\cdots$ repeat which is described by D. E. Knuth [23]. Statements between loop and repeat are iterated: when the while condition fails, the loop is exited. If the while were placed immediately following the loop, then the statement would be like a normal Algol while. In fact, Algorithm 1 might be implemented with a simpler construct like for


Diagram 4.
$c:=1$ until $N$ do $\cdots$ were it not for the need to test the control counter $c$ within the loop. The array $\mathrm{B}[N, c]$ is an index table which tells where the desired value of $\mathrm{P}[N]$ is after $\mathrm{P}[1], \cdots, \mathrm{P}[N-1]$ have been run through all permutations for the $c$ th time.

We still need to specify how to compute $\mathrm{B}[N, \mathrm{c}]$. For each value of $N$ we could specify any one of ( $N-1$ )! sequences in which to fill $\mathrm{P}[N]$, so there are a total of ( $N-1$ )! $(N-2)!(N-3)!\cdots 3!2!1!$ different tables $\mathrm{B}[N, c]$ which will cause Algorithm 1 to properly generate all $N$ ! permutations of $\mathrm{P}[1], \cdots, \mathrm{P}[N]$.

One possibility is to precompute $\mathrm{B}[N, c]$ by hand (since we know that $N$ is small), continuing as in the example above. If we adopt the rule that $\mathrm{P}[N]$ should be filled with elements in decreasing order of their original index, then the network in Diagram 4 tells us that B[5,c] should be $1,3,1,3$ for $c=1,2,3,4$. For $N=6$ we proceed in the same way: if we start with A B C D E F, then the first $N=5$ subnetwork leaves the elements in the order C D E B A F, so that $\mathrm{B}[6,1]$ must be 3 to get the E into $\mathrm{P}[6]$, leaving C D F B A E. The second $N=5$ subnetwork then leaves F B A D C E, so that $B[6,2]$ must be 4 to get the $D$ into $P[6]$, etc. Table 2 is the full table for $\mathrm{N} \leq 12$ generated this way; we could generate permutations with Algorithm 1 by storing these $N(N-1)$ indices.

There is no reason to insist that $\mathrm{P}[N]$ should be filled with elements in decreasing order. We could proceed as above to build a table which fills $\mathrm{P}[N]$ in any order we choose. One reason for doing so would be to try to avoid having to store the table: there are at least two known versions of

TABLE 2. Index Table $B[N, c]$ for Algorithm 1

|  | 2 | 1 |  |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 3 | 1 | 1 |  |  |  |  |  |  |  |  |  |
|  | 4 | 1 | 2 | 3 |  |  |  |  |  |  |  |  |
|  | 5 | 3 | 1 | 3 | 1 |  |  |  |  |  |  |  |
|  | 6 | 3 | 4 | 3 | 2 | 3 |  |  |  |  |  |  |
| $N$ | 7 | 5 | 3 | 1 | 5 | 3 | 1 |  |  |  |  |  |
|  | 8 | 5 | 2 | 7 | 2 | 1 | 2 | 3 |  |  |  |  |
|  | 9 | 7 | 1 | 5 | 5 | 3 | 3 | 7 | 1 |  |  |  |
|  | 10 | 7 | 8 | 1 | 6 | 5 | 4 | 9 | 2 | 3 |  |  |
|  | 11 | 9 | 7 | 5 | 3 | 1 | 9 | 7 | 5 | 3 | 1 |  |
|  | 12 | 9 | 6 | 3 | 10 | 9 | 4 | 3 | 8 | 9 | 2 | 3 |

this method in which the indices can be easily computed and it is not necessary to precompute the index table.

The first of these methods was one of the earliest permutation generation algorithms to be published, by M. B. Wells in 1960 [47]. As modified by J. Boothroyd in 1965 [1, 2], Wells' algorithm amounts to using

$$
\mathrm{B}[N, c]= \begin{cases}N-c & \text { if } N \text { is even and } c>2 \\ N-1 & \text { otherwise },\end{cases}
$$

or, in Algorithm 1, replacing $\mathrm{P}[\mathrm{B}[N, c]]:=: \mathrm{P}[N]$ by
if ( $N$ even) and ( $c>2$ )
then $\mathrm{P}[N]:=: \mathrm{P}[N-c]$
else $\mathrm{P}[N]$ :=: $\mathrm{P}[\mathrm{N}-1]$ endif
It is rather remarkable that such a simple method should work properly. Wells gives a complete formal proof in his paper, but many readers may be content to check the method for all practical values of $N$ by constructing the networks as shown in the example above. The complete networks for $N=2,3,4$ are shown in Fig. 2.

In a short paper that has gone virtually unnoticed, B.R. Heap [16] pointed out several of the ideas above and described a method even simpler than Wells'. (It is not clear whether Heap was influenced by Wells or Boothroyd, since he gives no references.) Heap's method is to use

$$
\mathrm{B}(N, c)= \begin{cases}1 & \text { if } N \text { is odd } \\ c & \text { if } N \text { is even },\end{cases}
$$

or, in Algorithm 1, to replace $\mathrm{P}[\mathrm{B}[N, \mathrm{c}]]:=: \mathrm{P}[N]$ by
if $N$ odd then $\mathrm{P}[N]$ :=: $\mathrm{P}[1]$ else $\mathrm{P}[N]:=: \mathrm{P}[c]$ endif
Heap gave no formal proof that his method works, but a proof similar to Wells' will show that the method works for all $N$. (The reader may find it instructive to verify that the method works for practical values of $N$ (as Heap did) by proceeding as we did when constructing the index table above.) Figure 3 shows that the networks for $N=2,3,4$ are the same as for Algorithm 1 with the precomputed index table, but that the network for $N=5$, shown in Diagram 5, differs. (The empty boxes denote the network for $N=4$ from Fig. 3.)


Figure 2. Wells' algorithm for $N=2,3,4$.

Neither Wells nor Heap gave recursive formulations of their methods, although Boothroyd [1] later gave a recursive implementation of his version of Wells' method. Although Wells and Heap undoubtedly arrived at their nonrecursive programs directly, it is instructive here to derive a nonrecursive version of Algorithm 1 by systematically removing the recursion.

The standard method for implementing a recursive procedure is to maintain a stack with the parameters and local variables for each invocation of the procedure. The simple structure of Algorithm 1 makes it more convenient to maintain an array $c[1], \cdots, c[N]$, where $c[i]$ is the value of $c$ for the invocation permutations $(i)$. Then by decrementing $i$ on a call and incrementing $i$ on return, we ensure that $c[i]$ always refers to the proper value of $c$. Since there is only one recursive call, transfer of control is implemented by jumping to the beginning on call and
jumping to the place following the call on return. The following program results directly when we remove the recursion and the loops from Algorithm 1:

```
\(t:=N\);
begin: c[l]:=1;
loop: if \(t>2\) then \(t:=t-1\); go to begin endif;
return: if \(\mathrm{c}[\nu] \geq \iota\) then go to extt endif;
    \(\mathrm{P}[\mathrm{B}[c[l]]]=: \mathrm{P}[l] ;\)
\(c[\imath 7:=c[\nu]+1\);
go to loop;
exit. if \(l<N\) then \(t:=\imath+1\); go to return endif;
```

This program can be simplified by combining the instructions at begin and loop into a single loop, and by replacing the single go to exit with the code at exit:

The program can be transformed further if we observe that, after $c[N], \cdots, c[2]$ are all set to 1 (this is the first thing that the program does), we can do the assignment $c[i]:=1$ before $t:=i+1$ rather than after $i=i-1$ without affecting the rest of the program. But this means that the loop does nothing but set $i$ to 2 and it can be eliminated (except for the initialization), as in this version:

```
\(t:=N\);
    loop: \(c[l \tau:=1\) while \(l>2 . v:=\imath-1\) repeat;
then if \(t<N\) then \(c[\imath]:=1 ; \imath:=\imath+1\);
                                    go to return endif;
else \(\mathrm{P}[\mathrm{B}[\mathrm{c}[\imath]]]:=. \mathrm{P}[\imath]\);
    \(c[\imath]:=c[l]+1\);
    \(t:=2\);
        go to return;
endif;
```

return if c $l \nu \geq 1$

Finally, since the two go to's in this program refer to the same label, they can be replaced with a single loop $\cdots$ repeat. The formulation
$i:=N$; loop: $c[\imath]:=1$ while $\imath>2: \imath:=\imath-1$ repeat;
loop:
if $c[\iota]<\iota$ then $\mathrm{P}[\mathrm{B}[c[\iota]]]:=: \mathrm{P}[\imath]$;
$c[\imath]:=c / \imath]+1, t:=2 ;$
else $c[\imath]:=1 ; \imath:=\imath+1$;
endif;
while $\iota \leq N$ repeat;


Figure 3. Heap's algorithm for $N=2,3,4$.
is attractive because it is symmetric: each time through the loop, either c $c i]$ is initialized and $i$ incremented, or $c[i]$ is incremented and $i$ initialized. (Note: in a sense, we have removed too many go to's, since now the program makes a redundant test $i$ $\leq N$ after setting $i:=2$ in the then clause. This can be avoided in assembly language, as shown in Section 3, or it could be handled with an "event variable" as described in [24].) We shall examine the structure of this program in detail later.

The programs above merely generate all permutations of $\mathrm{P}[1], \cdots, \mathrm{P}[N]$; in order to do anything useful, we need to process each permutation in some way. The processing might involve anything from simple counting to a complex simulation. Nor-


Diagram 5.
mally, this is done by turning the permutation generation program into a procedure which returns a new permutation each time it is called. A main program is then written to call this procedure $N$ ! times, and process each permutation. (In this form, the permutation generator can be kept as a library subprogram.) A more efficient way to proceed is to recognize that the permutation generation procedure is really the "main program" and that each permutation should be processed as it is generated. To indicate this clearly in our programs, we shall assume a macro called process which is to be invoked each time a new permutation is ready. In the nonrecursive version of Algorithm 1 above, if we put a call to process at the beginning and another call to process after the exchange statement, then process will be executed $N!$ times, once for each permutation. From now on, we will explicitly include such calls to process in all of our programs.

The same transformations that we applied to Algorithm 1 yield this nonrecursive version of Heap's method for generating and processing all permutations of $\mathrm{P}[1], \cdots, \mathrm{P}[N]:$

## Algorithm 2 (Heap)

$\imath:=N$; loop: $c[\iota]:=1$ while $\tau>2: u:=\iota-1$ repeat; process;
loop-
if $c / \downarrow<\imath$
then if $\imath$ odd then $k:=1$ else $k:=c[l]$ endif; $\mathrm{P}[\ell]:=: \mathrm{P}[k] ;$ $c[\imath]:=c[\imath]+1 ; \tau:=2$; process,
else $c\{\imath]:=1 ; \imath:=\imath+1$
endif;
while $1 \leq N$ repeat;
This can be a most efficient algorithm when implemented properly. In Section 3 we examine further improvements to this algorithm and its implementation.

## Adjacent Exchanges

Perhaps the most prominent permutation enumeration algorithm was formulated in 1962 by S. M. Johnson [20] and H. F. Trotter [45], apparently independently. They discovered that it was possible to generate all $N$ ! permutations of $N$ elements with $N!-1$ exchanges of adjacent elements.


The method is based on the natural idea that for every permutation of $N-1$ elements we can generate $N$ permutations of $N$ elements by inserting the new element into all possible positions. For example, for five elements, the first four exchange modules in the permutation network are as shown in Diagram 6. The next exchange is $P[1]:=: P[2]$, which produces a new permutation of the elements originally in $\mathrm{P}[2]$, $\mathrm{P}[3], \mathrm{P}[4], \mathrm{P}[5]$ (and which are now in $\mathrm{P}[1]$, $\mathrm{P}[2], \mathrm{P}[3], \mathrm{P}[4]$ ). Following this exchange, we bring A back in the other direction, as illustrated in Diagram 7. Now we exchange $\mathrm{P}[3]:=: \mathrm{P}[4]$ to produce the next permutation of the last four elements, and continue in this manner until all 4 ! permutations of the elements originally in $\mathrm{P}[2]$, P[3], P[4], P[5] have been generated. The network makes five new permutations of the five elements for each of these (by putting the element originally in $\mathrm{P}[1]$ in all possible positions), so that it generates a total of 5 ! permutations.
Generalizing the description in the last paragraph, we can inductively build the network for $N$ elements by taking the network for $N-1$ elements and inserting chains of $N-1$ exchange modules (to sweep the first element back and forth) in each space between exchange modules. The main complication is that the subnetwork for $N-1$ elements has to shift back and forth between the first $N-1$ lines and the last $N-1$ lines in between sweeps. Figure 4 shows the networks for $N=2,3,4$. The modules in boxes identify the subnetwork: if, in the network for $N$, we connect the output lines of one box to the input lines of the next, we get the network for $N-1$.
 Diagram 7.


Figure 4. Johnson-Trotter algorithm for $N=2$, 3, 4.

Continuing the example above, we get the full network for $N=5$ shown in Figure 5. By connecting the boxes in this network, we get the network for $N=4$.

To develop a program to exchange according to these networks, we could work down from a recursive formulation as in the preceding section, but instead we shall take a bottom-up approach. To begin, imagine that each exchange module is labelled with the number of the network in which it first appears. Thus, for $N=2$ the module would be numbered 2 ; for $N=3$ the five modules would be labelled 3323 ; for $N=4$ the 23 modules are numbered

$$
44434443444244434443444
$$

for $N=5$ we insert 5555 between the numbers above, etc. To write a program to generate this sequence, we keep a set of incrementing counters $c[i], 2 \leq i \leq N$, which are all initially 1 and which satisfy

Figure 5 Johnson-Trotter algorithm for $N=5$.
$1 \leq c[i] \leq i$. We find the highest index $i$ whose counter is not exhausted (not yet equal to $i$, output it, increment its counter, and reset the counters of the larger indices:

```
\(i:=1\); loop while \(i \leq N: i:=\iota+1 ; c[i]:=1\) repeat;
\(c[1]:=0\);
loop:
    \(i:=N\);
    loop while \(c[\imath]=i\) : \(c[i]:=1 ; i:=\iota-1\) repeat;
while \(l>1\) :
            comment exchange module is on level \(\boldsymbol{c}\);
    \(c[l]:=c[l]+1\)
repeat;
```

When $i$ becomes 1 , the process is completed - the statement $c[1]=0$ terminates the inner loop in this case. (Again, there are simple alternatives to this with "event variables" [24], or in assembly language.)

Now, suppose that we have a Boolean variable $d[N]$ which is true if the original $\mathrm{P}[1]$ is travelling from $\mathrm{P}[1]$ down to $\mathrm{P}[N]$ and false if it is travelling from $\mathrm{P}[N]$ up to $\mathrm{P}[1]$. Then, when $i=N$ we can replace the comment in the above program by

```
if \(d[N]\) then \(k:=c[N]\) else \(k:=N-c[N]\) endif;
\(\mathrm{P}[k]:=: \mathrm{P}[k+1]\);
```

This will take care of all of the exchanges on level $N$. Similarly, we can proceed by introducing a Boolean $d[N-1]$ for level $N-1$, etc., but we must cope with the fact that the elements originally in $\mathrm{P}[2], \cdots, \mathrm{P}[N]$ switch between those locations and $\mathrm{P}[1], \cdots, \mathrm{P}[N-1]$. This is handled by including an offset $x$ which is incremented by 1 each time a $d[i]$ switches from false to true. This leads to:

```
Algorithm 3 (Johnson-Trotter)
    \(t:=1\);
    loop while \(\imath<N: \imath:=\imath+1 ; c\{\imath\}:=1\);
        \(d[i]:=\) true; repeat;
    c[1]:=0;
    process;
    loop:
        \(i:=N ; x:=0\);
        loop while \(c\{\imath]=t\) :
            if not \(d[u]\) then \(x:=x+1\) endif;
            \(d[i]:=\operatorname{not} d[\iota] ; c[\iota]:=1 ; i .=\imath-1 ;\)
        repeat;
```

while $i>1$ :
if $d[\imath]$ then $k:=c[l]+x$
else $k:=\iota-c[\imath]+x$ endif;
$\mathrm{P}[k]:=: \mathrm{P}[k+1] ;$
process;
$c[i]:=c[i]+1 ;$
repeat;
Although Johnson did not present the algorithm in a programming language, he did give a very precise formulation from which the above program can be derived. Trotter gave an Algol formulation which is similar to Algorithm 3. We shall examine alternative methods of implementing this algorithm later.

An argument which is often advanced in favor of the Johnson-Trotter algorithm is that, since it always exchanges adjacent elements, the process procedure might be simpler for some applications. It might be possible to calculate the incremental effect of exchanging two elements rather than reprocessing the entire permutation. (This observation could also apply to Algorithms 1 and 2 , but the cases when they exchange nonadjacent elements would have to be handled differently.)

The Johnson-Trotter algorithm is often inefficiently formulated $[5,10,12]$ because it can be easily described in terms of the values of elements being permuted, rather than their positions. If $\mathrm{P}[1], \cdots, \mathrm{P}[N]$ are originally the integers $1, \cdots, N$, then we might try to avoid maintaining the offset $x$ by noting that each exchange simply involves the smallest integer whose count is not yet exhausted. Inefficient implementations involve actually searching for this smallest integer [5] or maintaining the inverse permutation in order to find it [10]. Both of these are far less efficient than the simple offset method of maintaining the indices of the elements to be exchanged given by Johnson and Trotter, as in Algorithm 3 .

## Factorial Counting

A careful reader may have become suspicious about similarities between Algo-
rithms 2 and 3. The similarities become striking when we consider an alternate implementation of the Johnson-Trotter method:

Algorithm $3 a$ (Alternate Johnson-Trotter)
$\imath=N ;$
loop: $c\{\iota]:=1 ; d[\iota\}:=$ true; while $\imath>1: \iota=\iota-1$ repeat;
process,
loop:
if $c\{l]<N+1-\imath$
then if $d[l]$ then $k .=c[l]+x$ else $k:=N+1-\imath-c[t]+x$ endif;
$\mathrm{P}[k]:=: \mathrm{P}[k+1]$;
process;
$c[\iota]:=c[t]+1 ; \tau:=1 ; x:=0 ;$
else if not $d[l]$ then $x:=x+1$ endif; $c[\imath]:=1 ; \imath:=\imath+1 ; d\lfloor\imath]:=$ not $d[\imath] ;$ endif;
while $t \leq N$ repeat;
This program is the result of two simple transformations on Algorithm 3. First, change $i$ to $N+1-l$ everywhere and redefine the $c$ and $d$ arrays so that $c[N+1-l]$, $d[N+1-i]$ in Algorithm 3 are the same as $c[i], d[i]$ in Algorithm 3a. (Thus a reference to $c[i]$ in Algorithm 3 becomes $c[N+1-i]$ when $i$ is changed to $N+1-i$, which becomes $c[i]$ in Algorithm 3a.) Second, rearrange the control structure around a single outer loop. The condition $c[i]<N+1-i$ in Algorithm 3a is equivalent to the condition $c[i]<i$ in Algorithm 3 , and both programs perform the exchange and process the permutation in this case. When the counter is exhausted ( $c[i]=N+1-i$ in Algorithm 3a; $c[i]=i$ in Algorithm 3), both programs fix the offset, reset the counter, switch the direction, and move up a level.

If we ignore statements involving $\mathrm{P}, k$ and $d$, we find that this version of the Johnson-Trotter algorithm is identical to Heap's method, except that Algorithm 3a compares $c[i]$ with $N+1-i$ and Algorithm 2 compares it with $i$. (Notice that Algorithm 2 still works properly if in both its occurrences 2 is replaced by 1 .)

To appreciate this similarity more fully, let us consider the problem of writing a program to generate all $N$-digit decimal numbers: to "count" from 0 to $99 \cdots 9=$ $10^{N}-1$. The algorithm that we learn in grade school is to increment the right-most
digit which is not 9 and change all the nines to its right to zeros. If the digits are stored in reverse order in the array $c[N], c[N-1], \cdots, c[2], c[1]$ (according to the way in which we customarily write numbers) we get the program

$$
\begin{aligned}
& \imath:=N, \text { loop } c l \imath j:=0 \text { while } \imath>1 \quad \imath:=\imath-1 \text { repeat; } \\
& \text { loop: } \\
& \quad \text { if } c[\imath]<9 \text { then } c[\imath]:=c[\imath]+1 ; \imath=1 \\
& \quad \text { endife } c[\imath]:=0 ; \imath=l+1
\end{aligned}
$$

From this program, we see that our permutation generation algorithms are controlled by this simple counting process, but in a mixed-radix number system. Where in ordinary counting the digits satisfy $0 \leq$ $c[i] \leq 9$, in Algorithm 2 they satisfy $1 \leq$ $c[i] \leq i$ and in Algorithm 3a they satisfy 1 $\leq c[i] \leq N-i+1$. Figure 6 shows the values of $c[1], \cdots, c[N]$ when process is encountered in Algorithms 2 and 3a for $N=$ 2,3,4.

Virtually all of the permutation generation algorithms that have been proposed are based on such "factorial counting" schemes. Although they appear in the literature in a variety of disguises, they all have the same control structure as the elementary counting program above. We have called methods like Algorithm 2 recursive because they generate all sequences of $c[1], \cdots, c[i-1]$ in-between increments of $c[i]$ for all $i$; we shall call methods like Algorithm 3 iterative because they iterate $c[i]$ through all its values inbetween increments of $c[i+1], \cdots, c[N]$.

## Loopless Algorithms

An idea that has attracted a good deal of attention recently is that the JohnsonTrotter algorithm might be improved by removing the inner loop from Algorithm 3. This idea was introduced by G. Ehrlich [10, 11], and the implementation was refined by N. Dershowitz [5]. The method is also described in some detail by S. Even [12].

Ehrlich's original implementation was complex, but it is based on a few standard programming techniques. The inner loop

|  | $N=4$ |  | $\mathrm{v}=4$ |
| :---: | :---: | :---: | :---: |
|  | 1111 |  | 1111 |
|  | 1121 |  | 1112 |
|  | - 1211 |  | 1113 |
|  | 1221 |  | 1114 |
|  | 1311 |  | 1121 |
|  | 1321 |  | 1122 |
| $N=2$ | 2111 | $N=2$ | 1123 |
| 11 |  | 11 |  |
|  | 2121 |  | 1124 |
| 21 | 2211 | 12 | 1131 |
|  | 2221 |  | 1132 |
|  | 23111 |  | 1133 |
| $N=3$ | 2321 | $N \pm 3$ | 1134 |
| 111 | 3111 | 111 | 1211 |
| 121 | 3121 | 112 | 1212 |
| 211 | 3211 | 113 | 1213 |
| 221 | 3221 | 121 | 1214 |
| 311 | 3311 | 122 | 1221 |
| 321 | 3321 | 123 | 1222 |
|  | 4111 |  | 1223 |
|  | 4121 |  | 1224 |
|  | 4211 |  | 1231 |
|  | 4221 |  | 1232 |
|  | $\begin{array}{lllll}4 & 1 & 1\end{array}$ |  | 1233 |
|  | 4321 |  | 1234 |
|  |  |  |  |

Figure 6. Factorial counting: $c[N], \cdots, c[1]$. rithm 3 a (iterative).
in Algorithm 3 has three main purposes: to find the highest index whose counter is not exhausted, to reset the counters at the larger indices, and to compute the offset $x$. The first purpose can be served by maintaining an array $s[i]$ which tells what the next value of $i$ should be when $c[i]$ is exhausted: normally $s[i]=i-1$, but when $c[i]$ reaches its limit we set $s[i+1]$ to $s[i]$. To reset the other counters, we proceed as we did when removing the recursion from Algorithm 1 and reset them just after they are incremented, rather than waiting until they are needed. Finally, rather than computing a "global" offset $x$, we can maintain an array $x[i]$ giving the current offset at level $i$ : when $d[i]$ switches from false to true, we increment $x[s[i]]$. These changes allow us to replace the inner "loop...repeat" in Algorithm 3 by an "if $\cdots$ endif".

Algorithm $3 b$ (Loopless Johnson-Trotter) $\iota:=0$;
loop while $\imath<N: i:=t+1 ; c[\imath]:=1 ; d[\imath]:=$ true; $s[i]:=\imath-1 ; x[i]:=0$ repeat;
process;
loop:
$s[N+1] .=N ; x[\imath]:=0 ;$
if $c[l]=i$ then
if not $d[i]$
then $x\lceil s[\imath]]:=x[s[\imath]]+1$; endif;
$d[\iota]:=\operatorname{not} d[\iota] ; c[i]:=1$;
$s[\imath+1]:=s[i] ; s[i]:=\imath-1 ;$
endif;
$\iota:=s[N+1] ;$
while $l>1$ :
if $d[l]$ then $k:=c[l]+x[l]$
else $k:=l-c[l]+x[i]$ endif;
$\mathrm{P}[k]:=: \mathrm{P}[k+1]$;
process;
$c[\iota]:=c[\iota]+1 ;$
repeat;
This algorithm differs from those described in [10, 11, 12], which are based on
the less efficient implementations of the Johnson-Trotter algorithm mentioned above. The loopless formulation is purported to be an improvement because each iteration of the main loop is guaranteed to produce a new permutation in a fixed number of steps.
However, when organized in this form, the unfortunate fact becomes apparent that the loopless Algorithm 3b is slower than the normal Algorithm 3. Loopfree implementation is not an improvement at all! This can be shown with very little analysis because of the similar structure of the algorithms. If, for example, we were to count the total number of times che statement $c[i]:=1$ is executed when each algorithm generates all $N$ ! permutations, we would find that the answer would be exactly the same for the two algorithms. The loopless algorithm does not eliminate any such assignments; it just rearranges their order of execution. But this simple fact means that Algorithm 3b must be slower than Algorithm 3, because it not only has to execute all of the same instructions the same number of times, but it also suffers the overhead of maintaining the $x$ and $s$ arrays.

We have become accustomed to the idea that it is undesirable to have programs with loops that could iterate $N$ times, but this is simply not the case with the John-son-Trotter method. In fact, the loop iterates $N$ times only once out of the $N!$ times that it is executed. Most often ( $N-1$ out of every $N$ times) it iterates only once. If $N$ were very large it would be conceivable that the very few occasions that the loop iterates many times might be inconvenient, but since we know that $N$ is small, there seems to be no advantage whatsoever to the loopless algorithm.
Ehrlich [10] found his algorithm to run "twice as fast" as competing algorithms, but this is apparently due entirely to a simple coding technique (described in Section 3) which he applied to his algorithm and not to the others.

## Another Iterative Method

In 1976, F. M. Ives [19] published an ex-change-based method like the JohnsonTrotter method which does represent an improvement. For this method, we build

up the network for $N$ elements from the network for $N-2$ elements. We begin in the same way as in the Johnson-Trotter method. For $N=5$, the first four exchanges are as shown in Diagram 6. But now the next exchange is $P[1]:=: P[5]$, which not only produces a new permutation of $\mathrm{P}[1], \cdots, \mathrm{P}[4]$, but also puts $\mathrm{P}[5]$ back into its original position. We can perform exactly these five exchanges four more times, until, as shown in Diagram 8, we get back to the original configuration.

At this point, $\mathrm{P}[1], \cdots, \mathrm{P}[4]$ have been rotated through four permutations, so that we have taken care of the case $N=4$. If we (inductively) permute three of these elements (Ives suggests the middle three) then the 20 exchanges above will give us 20 new permutations, and so forth. (We shall later see a method which makes exclusive use of this idea that all permutations of $N$ elements can be generated by rotating and then generating all permutations of $N-1$ elements.) Figure 7 shows the networks for $N=2,3,4$; the full network for $N=5$ is shown in Fig. 8. As before, if we connect the boxes in the network for $N$, we get the network for $N-2$. Note that the exchanges immediately preceding the boxes are redundant in that they do not produce new permutations. (The redundant permutations are identified by parentheses in Fig. 7.) However, there are relatively few of these and they are a small price to pay for the lower overhead incurred by this method.
In the example above, we knew that it was time to drop down a level and permute the middle three elements when all of the original elements (but specifically $\mathrm{P}[1]$ and $P[5])$ were back in position. If the elements being permuted are all distinct, we can test for this condition by intially saving the values of $\mathrm{P}[1], \cdots, \mathrm{P}[N]$ in another array $\mathrm{Q}[1], \cdots, \mathrm{Q}[N]$ :
Algorithm 4 (Ives)
$t:=N$;
loop: $c f l]=t ; \mathrm{Q}[t],=\mathrm{P}[t]$, while $\imath<1 . \imath=t-1$ repeat; process,

```
loop.
    if \(c[u]<N+1-l\)
    then \(\mathrm{P}[c[t]]:=: \mathrm{P}[c[\imath]+1]\)
        \(c \mid \imath t=c[\imath]+1 ; u:=1 ;\)
        process;
    else \(\mathrm{P}[\imath] .=: \mathrm{P}[N+1-\imath]\);
    \(c[\imath]:=\imath ;\)
        if \(\mathrm{P}[N+1-\imath]=\mathrm{Q}[N+1-\imath]\) then \(t:=\imath+1\)
            else \(t:=1\);
                            process
        endif,
    endif,
while \(\imath<N+1-\imath\) repeat,
```

This program is very similar to Algorithms 2 and 3a, but it doesn't fall immediately into the "factorial counting" schemes of these programs, because only half of the


Figure 7 Ives' algorithm for $N=2,3,4$
counters are used. However, we can use counters rather than the test $\mathbf{P}[N+1-i]=$ $\mathrm{Q}[N+1-i]$, since this test always succeeds after exactly $\imath-1$ sweeps. We are immediately led to the implementation:

```
Algorithm 4a (Alternate Ives)
    \(\imath=N\); loop: \(c / \imath\) ]: \(\mathbf{1}\); while \(\iota>1: t:=\imath-1\) repeat,
    process;
    loop
        if \(c[l]<N+1-\imath\)
        then if \(\iota\) odd then \(\mathrm{P}[c / \imath]]:=: \mathrm{P}[c[\imath]+1]\)
                            else \(\mathrm{P}[\imath]:=: \mathrm{P}[N+1-\imath]\) endif,
            \(c[\nu]=c[\imath]+1, t:=1 ;\)
        process;
        else \(c[\iota]:=1 ; \imath=\imath+1\);
        endif,
while \(\imath \leq N\) repeat;
```

This method does not require that the elements being permuted be distinct, but it is slightly less efficient than Algorithm 4 because more counting has to be done.

Ives' algorithm is more efficient than the Johnson-Trotter method (compare Algorithm 4a with Algorithm 3a) since it does not have to maintain the array $d$ or offset $x$. The alternate implementation bears a remarkable resemblance to Heap's method (Algorithm 2). Both of these algorithms do little more than factorial counting. We shall compare them in Section 3.

## 2. OTHER TYPES OF ALGORITHMS

In this section we consider a variety of algorithms which are not based on simple exchanges between elements. These algorithms generally take longer to produce all permutations than the best of the methods already described, but they are worthy of study for several reasons. For example, in some situations it may not be necessary to generate all permutations, but only some "random" ones. Other algorithms may be of practical interest because they are based on elementary operations which could be as efficient as exchanges on some computers. Also, we consider algorithms that generate the permutations in a particular order which is of interest. All of the algorithms can be cast in terms of the basic


Figure 8 Ives' algorithm for $N=5$
"factorial counting" control structure described above.

## Nested Cycling

As we saw when we examined Ives' algorithm, $N$ different permutations of $\mathrm{P}[1], \cdots, \mathrm{P}[N]$ can be obtained by rotating the array. In this section, we examine permutation generation methods which are based solely on this operation. We assume that we have a primitive operation
rotate(i)
which does a cyclic left-rotation of the elements $\mathrm{P}[1], \cdots, \mathrm{P}[i]$. In other words, rotate( $(i)$ is equivalent to

```
\(t:=\mathrm{P}[1] ; k:=2\);
loop while \(k \leq t: \mathrm{P}[k-1]\) := \(\mathrm{P}[k]\) repeat;
\(\mathrm{P}[i]:=t\);
```

The network notation for rotate(5) is given by Diagram 9 . This type of operation can be performed very efficiently on some computers. Of course, it will generally not be more efficient than an exchange, since rotate( 2 ) is equivalent to $\mathrm{P}[1]:=: \mathrm{P}[2]$.


Diagram 9.
The most straightforward way to make use of such an operation is a direct recursive implementation like Algorithm 1:

```
procedure permutations (N);
begin c:=1;
        loop:
            if N>2 then permutations(N-1)
            endif;
            rotate(N);
        while c<N
            process;
            c:=c+1
        repeat;
end;
```

When the recursion is removed from this program in the way that removed the recursion from Algorithm 1, we get an old algorithm which was discovered by C. Tompkins and L. J. Paige in 1956 [44]:

[^2]loop:
rotate(2)
if $c[u]<l$ then $c[u]:=c[i]+1 ; i:=2$;
process;
else $c[\iota]:=1 ; \imath:=\imath+1$
endif;
while $t \leq N$ repeat;
This is nothing more than a simple counting program with rotation added. The rotation networks and permutation sequences generated by this algorithm are given in Fig. 9. As in Fig. 7, the parenthesized permutations are redundant permutations produced in the course of the computation, which are not passed through the process macro. The algorithm is not as inefficient as it may seem, because most of the rotations are short, but it clearly will not compete with the algorithms in Section 1 . This method apparently represents the earliest attempt to get a real computer to generate permutations as quickly as possible. An Algol implementation was given by Peck and Schrack [33].

An interesting feature of the TompkinsPaige method, and the reason it generates so many redundant sequences, is that the recursive procedure restores the permutation to the order it had upon entry. Programs that work this way are called backtrack programs [26, 27, 46, 48]. We can easily apply the same idea to exchange methods like Algorithm 1. For example:

```
procedure permutatoons(N);
        begin c}=1
            loop:
            P[N]:=:P[c];
            if N>2 then permutatoons(N-1)
                        else process endif;
            P[c]:=:P[N];
                while}c<N\mathrm{ :
                    c:=c+1
            repeat;
        end;
```

A procedure like this was given by C. T. Fike [13], who also gave a nonrecursive version [37] which is similar to a program developed independently by S. Pleszczyński [35]. These programs are clearly less efficient than the methods of Section 1, which have the same control structure but require many fewer exchanges.

Tompkins was careful to point out that it is often possible easily to achieve great savings with this type of procedure. Often,


Figure 9. Tompkins-Paige algorithm for $N=2$, 3,4.
process involves selecting a small set of permutations satisfying some simple criteria. In certain cases, once a permutation is found not to satisfy the criteria, then the
permutations generated by permuting its initial elements will not satisfy the criteria either, so the exchanges and the call permutations $(N-1)$ can be skipped. Thus large numbers of permutations never need be generated. A good backtracking program will eliminate nearly all of the processing (see [44] for a good example), and such methods are of great importance in many practical applications.

Cycling is a powerful operation, and we should expect to find some other methods which use it to advantage. In fact, Tompkins' original paper [44] gives a general proof from which several methods can be constructed. Remarkably, we can take Algorithm 5 and switch to the counter system upon which the iterative algorithms in Section 2 were based:

$$
\begin{aligned}
& \imath:=N \text {; loop: } c\{\imath\}:=1 \text { while } \iota>1 . t:=t-1 \text { repeat; } \\
& \text { process; } \\
& \text { loop: } \\
& \text { rotate }(N+1-\imath) \text {; } \\
& \text { if } c[l]<N+1-\imath \text { then } c[l]:=c[l]+1 ; \imath:=1 \text {; } \\
& \text { process; } \\
& \text { else } c[t]:=1 ; \imath:=\imath+1
\end{aligned}
$$

endif; while $t \leq N$ repeat,

Although fewer redundant permutations are generated, longer rotations are involved, and this method is less efficient than Algorithm 5. However, this method does lend itself to a significant simplification, similar to the one Ives used. The condition $c[i]=N+1-\imath$ merely indicates that the elements in $\mathrm{P}[1], \mathrm{P}[2], \cdots$, $\mathrm{P}[N+1-i]$ have undergone a full rotation - that is, $\mathrm{P}[N+1-\imath]$ is back in its original position in the array. This means that if we initially set $\mathrm{Q}[i]=\mathrm{P}[i]$ for $1 \leq i \leq N$, then $c[i]=N+1-i$ is equivalent to $\mathrm{P}[N+1-i]=\mathrm{Q}[N+1-i]$. But we have now removed the only test on $c[i]$, and now it is not necessary to maintain the counter array at all! Making this simplification, and changing $\iota$ to $N+1-i$, we have an algorithm proposed by G. Langdon in 1967 [25], shown in Fig. 10.

## Algorithm 6 (Langdon)

[^3]
## loop:

## rotate( $\imath$ );

if $\mathrm{P}[t]=\mathrm{Q}[\iota]$ then $t:=N$ else $\tau=\iota-1$ endif; process;
while $t \geq 1$ repeat;
This is definitely the most simply expressed of our permutation generation algorithms. If $\mathrm{P}[1], \cdots, \mathrm{P}[N]$ are initially $1, \cdots, N$, then we can eliminate the initialization loop and replace $\mathrm{Q}[i]$ by $i$ in the main loop. Unfortunately, this algorithm runs slowly on most computers-only


Figure 10. Langdon's algorithm for $N=2,3,4$.
when very fast rotation is available is it the method of choice. We shall examine the implementation of this algorithm in detail in Section 3.

## Lexicographic Algorithms

A particular ordering of the $N$ ! permutation of $N$ elements which is of interest is "lexicographic", or alphabetical, ordering. The lexicographic ordering of the 24 permutations of A B C D is shown in Fig. 11a. "Reverse lexicographic" ordering, the result of reading the lexicographic sequence backwards and the permutations from right to left, is also of some interest. Fig. 11b shows the reverse lexicographic ordering of the 24 permutations of A B C D.

The natural definition of these orderings has meant that many algorithms have been proposed for lexicographical permutation generation. Such algorithms are inherently less efficient than the algorithms of Section 1, because they must often use more than one exchange to pass from one permutation to the next in the sequence (e.g., to pass from A C D B to A D B C in Fig. 11a). The main practical reason that has been advanced in favor of lexicographic generation is that, in reverse order, all permutations of $\mathrm{P}[1], \cdots, \mathrm{P}[N-1]$ are generated before $\mathrm{P}[N]$ is moved. As with backtracking, a processing program could, for example, skip ( $N-1$ )! permutations in some instances. However, this property is shared by the recursive algorithms of Section 1 -in fact, the general structure of Algorithm 1 allows $\mathrm{P}[N]$ to be filled in any arbitrary order, which could be even more of an advantage than lexicographic ordering in some instances.

Nevertheless, lexicographic generation is an interesting problem for which we are by now well prepared. We shall begin by assuming that $\mathrm{P}[1], \cdots \mathrm{P}[N]$ are distinct. Otherwise, the problem is quite different, since it is usual to produce a lexicographic listing with no duplicates (see [3, 9, 39].

We shall find it convenient to make use of another primitive operation, reverse( $i$ ). This operation inverts the order of the elements in $\mathrm{P}[1], \cdots \mathrm{P}[i]$; thus, reverse( $i$ ) is equivalent to

| $A B C D$ | $A B C D$ |
| :---: | :---: |
| A B D C | B A C D |
| $A C B D$ | $A C B D$ |
| A CDB | $C A B D$ |
| A D B C | $B C A D$ |
| A DCB | $C \mathrm{BAD}$ |
| $B \mathrm{ACD}$ | A B DC |
| $B A D C$ | BADC |
| $B C A D$ | A D B C |
| $B C D A$ | DABC |
| $B \mathrm{DAC}$ | B DAC |
| $B \mathrm{DCA}$ | D BAC |
| C A B D | $A C D B$ |
| $C A D B$ | CA D B |
| $C B A D$ | A D C B |
| $C \mathrm{BDA}$ | D A C B |
| C D A B | C D A b |
| C D B A | D $C$ A B |
| D A B C | $B C D A$ |
| D ACB | CBDA |
| D BAC | B D C A |
| D B C A | D B C A |
| DCAB | C DBA |
| D C B A | D C B A |

(a)

ABCD
B A C D
ACBD
$B C A D$
C BAD
ABDC
$A D B C$
DABC
B DAC
$A C D B$
CADB
A D C B

CDAB
D C A B
$B C D A$
CBDA

DECA
C DBA
D C B A
(b)

Figure 11. Lexicographic ordering of A B C D.
(a) In natural order. (b) In reverse order.

## $\imath:=1$;

loop while $i<N+1-l$ :

$$
\mathrm{P}[\imath]:=: \mathrm{P}[N+1-\imath] ; z:=\imath+1 \text { repeat; }
$$

This operation will not be particularly efficient on most real computers, unless special hardware is available. However, it seems to be inherent in lexicographic generation.

The first algorithm that we shall consider is based on the idea of producing each permutation from its lexicographic predecessor. Hall and Knuth [15] found that the method has been rediscovered many times since being published in 1812 by Fischer and Krause [14]. The ideas involved first appear in the modern literature in a rudimentary form in an algorithm by G. Schrack and M. Shimrat [40]. A full formulation was given by M. Shen in 1962 [41, 42], and Phillips [34] gives an "optimized" implementation. (Dijkstra [6] cites the problem as an example to illustrate a
"dramatic improvement in the state of the art" of computer programming since the algorithm is easily expressed in modern languages; the fact that the method is over 160 years old is perhaps a more sobering comment on the state of the art of computer science.)

The direct lexicographic successor of the permutation

BACFHGDE
is clearly

BACFHGED,

but what is the successor of this permutation? After some study, we see that H G E D are in their lexicographically highest position, so the next permutation must begin as B A C G $\cdots$. The answer is the lexicographically lowest permutation that begins in this way, or

BACGDEFH.

Similarly, the direct successor of
HFEDGCAB
in reverse lexicographic order is

> D E G H F C A B,
and its successor is

## EDGHFCAB.

The algorithm to generate permutations in reverse lexicographic order can be clearly understood from this example. We first scan from left to right to find the first $i$ such that $\mathrm{P}[l]>\mathrm{P}[i-1]$. If there is no such $i$, then the elements are in reverse order and we terminate the algorithm since there is no successor. (For efficiency, it is best to make $\mathrm{P}[N+1]$ larger than all the other elements-written $\mathrm{P}[N+1]=$ $\infty-$ and terminate when $i=N+1$.) Otherwise, we exchange $P[i]$ with the next-lowest element among $\mathrm{P}[1], \cdots, \mathrm{P}[i-1]$ and then reverse $P[1], \cdots, P[i-1]$. We have

```
Algorithm 7 (Fischer-Krause)
    \(\mathrm{P}[N+1]=\infty\);
    process;
    loop.
        \(\imath:=2\); loop while \(\mathrm{P}[\imath]<\mathrm{P}[\imath-1]: \imath:=\imath+1\) repeat;
    while \(t<N\);
    \(J:=1\); loop while \(\mathrm{P}[J]>\mathrm{P}[i]: J:=\jmath+1\) repeat;
    \(\mathrm{P}[\) ] \(]=: \mathrm{P}[\mathrm{]}]\);
```

reverse $(\imath-1)$;
process;
repeat;
Like the Tompkins-Paige algorithm, this algorithm is not as inefficient as it seems, since it is most often scanning and reversing short strings.

This seems to be the first example we have seen of an algorithm which does not rely on "factorial counting" for its control structure. However, the control structure here is overly complex; indeed, factorial counters are precisely what is needed to eliminate the inner loops in Algorithm 7.
A more efficient algorithm can be derived, as we have done several times before, from a recursive formulation. A simple recursive procedure to generate the permutations of $\mathrm{P}[1], \cdots, \mathrm{P}[N]$ in reverse lexicographic order can be quickly devised:

```
procedure lexperms(N);
    begin c:=1;
            loop:
                    if N>2
                        then lexperms(N-1) endif;
                while}c<N\mathrm{ :
                    P[N]:=:P[c];
                    reverse(N-1);
                    c:=c+1;
                repeat;
        end;
```

Removing the recursion precisely as we did for Algorithm 1, we get

```
Algorithm 8 (Ord-Smith)
    \(\imath:=N\); loop \(c[\imath]\). \(=1\); while \(\imath>2: \imath:=\imath-1\) repeat;
    process;
    loop:
        if \(c(l]<l\) then \(\mathrm{P}[\imath]:=\cdot \mathrm{P}[c f l]\), reverse \((\imath-1)\),
        \(c[\imath]:=c[\imath]+1 ; \imath:=2 ;\)
                        process;
                else \(c\{l\rangle:=1 ; \imath:=\imath+1\)
        endif;
    while \(t \leq N\) repeat;
```

This algorithm was first presented by R. J. Ord-Smith in 1967 [32]. We would not expect a priori to have a lexicographic algorithm so similar to the normal algorithms, but the recursive formulation makes it obvious.
Ord-Smith also developed [31] a "pseudo-lexicographic" algorithm which consists of replacing $\mathrm{P}[\imath]:=: \mathrm{P}[c[i]]$; reverse( $(i-1)$; by reverse( $i$ ) in Algorithm 8.

There seem to be no advantages to this method over methods like Algorithm 1. Howell [17, 18] gives a lexicographic method based on treating $\mathrm{P}[1], \cdots, \mathrm{P}[N]$ as a base- $N$ number, counting in base $N$, and rejecting numbers whose digits are not distinct. This method is clearly very slow.

## Random Permutations

If $N$ is so large that we could never hope to generate all permutations of $N$ elements, it is of interest to study methods for generating "random" permutations of $N$ elements. This is normally done by establishing some one-to-one correspondence between a permutation and a random number between 0 and $N!-1$. (A full treatment of pseudorandom number generation by computer may be found in [22].)

First, we notice that each number between 0 and $N!-1$ can be represented in a mixed radix system to correspond to an array $c[N], c[N-1], \cdots, c[2]$ with $0 \leq c[i] \leq$ ${ }_{i-1}$ for $2 \leq i \leq N$. For example, 1000 corresponds to 121220 since $1000=6$ ! $+2 \cdot 5!+4!+2 \cdot 3!+2 \cdot 2!$. For $0 \leq n<N!$, we have $n=c[2] \cdot 1!+c[3] \cdot 2!+\cdots+$ $c[N] \cdot(N-1)!$. This correspondence is easily established through standard radix conversion algorithms [22, 27, 47]. Alternatively, we could fill the array by putting a "random" number between 0 and $i-1$ in $c[i]$ for $2 \leq i \leq N$.

Such arrays $c[N], c[N-1], \cdots, c[2]$ can clearly be generated by the factorial counting procedure discussed in Section 1, so that there is an implicit correspondence between such arrays and permutations of $12 \cdots N$. The algorithms that we shall examine now are based on more explicit correspondences.
The first correspondence has been attributed to M. Hall, Jr., in 1956 [15], although it may be even older. In this correspondence, $c[i]$ is defined to be the number of elements to the left of $i$ which are smaller than it. Given an array, the following example shows how to construct the corresponding permutation. To find the permutation of $12 \cdots 7$ corresponding to
we begin by writing down the first element, 1 . Since $c[2]=0,2$ must precede 1 , or

$$
21
$$

Similarly, $c[3]=2$ means that 3 must be preceded by both 2 and 1 :

$$
213
$$

Proceeding in this manner, we build up the entire permutation as follows:

$$
\begin{array}{lllllll}
2 & 1 & 4 & 3 & & \\
2 & 5 & 1 & 4 & 3 & \\
2 & 5 & 6 & 1 & 4 & 3 \\
2 & 7 & 5 & 6 & 1 & 4 & 3
\end{array}
$$

In general, if we assume that $c[1]=0$, we can construct the permutation $\mathrm{P}[1], \cdots$, $\mathrm{P}[N]$ with the program

```
\(r=1 ;\)
loop.
    \(j:=2\);
    loop while \(J>c[\tau]+\mathbf{1}: \mathrm{P} \cup] .=\mathrm{P} \cup-1]\); repeat;
    \(\mathrm{P}[c[\imath]+1]:=\imath\);
    \(i=\imath+1\);
while \(t \leq N\) repeat;
```

This program is not particularly efficient, but the correspondence is of theoretic interest. In a permutation $\mathrm{P}[1], \cdots, \mathrm{P}[N]$ of $1, \cdots, N$, a pair $(i, j)$ such that $i<j$ and $\mathrm{P}[i]>\mathrm{P}[]$ is called an inversion. The counter array in Hall's correspondence counts the number of inversions in the corresponding permutation and is called an inversion table. Inversion tables are helpful combinatorial devices in the study of several sorting algorithms (see [23]). Another example of the relationship between inversion tables and permutation generation can be found in [7], where Dijkstra reinvents the Johnson-Trotter method using inversion tables.
D. H. Lehmer [26, 27] describes another correspondence that was defined by D . N. Lehmer as long ago as 1906. To find the permutation corresponding to 121220 , we first increment each element by 1 to get

$$
232331 .
$$

Now, we write down 1 as before, and for $i=2, \cdots, N$, we increment all numbers which are $\geq c[i]$ by 1 and write $c[i]$ to the left. In this way we generate

$$
\begin{array}{r} 
\\
\\
\\
\\
\\
\\
\\
\\
\\
\\
\\
\\
\\
\\
2
\end{array}
$$

so that the permutation 2436715 corresponds to 1000 . In fact, 2436715 is the 1000th permutation of $1, \cdots, 7$ in lexicographic order: there are 6! permutations before it which begin with 1 , then $2 \cdot 5$ ! which begin 21 or 23 , then 4 ! which begin 241 , then $2 \cdot 3$ ! which begin 2431 or 2435 , and then $2 \cdot 2$ ! which begin 24361 or 24365 .

A much simpler method than the above two was apparently first published by $R$. Durstenfeld [8]. (See also [22], p. 125). We notice that $\mathrm{P}[i]$ has $i-1$ elements preceding it, so we can use the $c$ array as follows:

```
\(\imath:=N\);
loop while \(\imath \geq 2: \mathrm{P}[\tau]:=: \mathrm{P}[c[\tau]+1] ; \iota:=\imath-1\) repeat;
```

If we take $\mathrm{P}[1], \cdots, \mathrm{P}[\mathrm{N}]$ to be initially $1, \cdots, N$, then the array 121220 corresponds to the permutation 5146732 . This method involves only one scan through the array and is clearly more efficient than the above two methods.

We could easily construct a program to generate all permutations of $12 \cdots N$ by embedding one of the methods described above in a factorial counting control structure as defined in Section 1. Such a program would clearly be much slower than the exchange methods described above, because it must build the entire array $\mathrm{P}[1], \cdots, \mathrm{P}[N]$ where they do only a simple exchange. Coveyou and Sullivan [4] give an algorithm that works this way. Another method is given by Robinson [37].

## 3. IMPLEMENTATION AND ANALYSIS

The analysis involved in comparing a number of computer algorithms to perform the same task can be very complex. It is often necessary not only to look carefully at how the algorithms will be implemented on real computers, but also to carry out some complex mathematical analysis. Fortunately, these factors pres-
ent less difficulty than usual in the case of permutation generation. First, since all of the algorithms have the same control structure, comparisons between many of them are immediate, and we need only examine a few in detail. Second, the analysis involved in determining the total running time of the algorithms on real computers (by counting the total number of times each instruction is executed) is not difficult, because of the simple counting algorithms upon which the programs are based.

If we imagine that we have an important application where all $N$ ! permutations must be generated as fast as possible, it is easy to see that the programs must be carefully implemented. For example, if we are generating, say, every permutation of 12 elements, then every extraneous instruction in the inner loop of the program will make it run at least 8 minutes longer on most computers (see Table 1).

Evidently, from the discussion in Section 1, Heap's method (Algorithm 2) is the fastest of the recursive exchange algorithms examined, and Ives' method (Algorithm 4) is the fastest of the iterative exchange algorithms. All of the algorithms in Section 2 are clearly slower than these two, except possibly for Langdon's method (Algorithm 6) which may be competitive on machines offering a fast rotation capability. In order to draw conclusions comparing these three algorithms, we shall consider in detail how they can be implemented in assembly language on real computers, and we shall analyze exactly how long they can be expected to run.

As we have done with the high-level language, we shall use a mythical assembly language from which programs on real computers can be easily implemented. (Readers unfamiliar with assembly language should consult [21].) We shall use load (LD), store (ST), add (ADD), subtract (SUB), and compare (CMP) instructions which have the general form

## LABEL OPCODE REGISTER, OPERAND (optional)

The first operand will always be a symbolic register name, and the second operand may be a value, a symbolic register
name, or an indexed memory reference. For example, ADD I, 1 means "increment Register I by 1"; ADD I,J means "add the contents of Register J to Register I"; and ADD I,C(J) means "add to Register I the contents of the memory location whose address is found by adding the contents of Register J to C ". In addition, we shall use control transfer instructions of the form

## OPCODE LABEL

namely JMP (unconditional transfer); JL, JLE, JE, JGE, JG (conditional transfer according as whether the first operand in the last CMP instruction was $<, \leq,=, \geq,>$ than the second); and CALL (subroutine call). Other conditional jump instructions are of the form

OPCODE REGISTER, LABEL
namely $J N, J Z, J P$ (transfer if the specified register is negative, zero, positive). Most machines have capabilities similar to these, and readers should have no difficulty translating the programs given here to particular assembly languages.

Much of our effort will be directed towards what is commonly called code optimization: developing assembly language implementations which are as efficient as possible. This is, of course, a misnomer: while we can usually improve programs, we can rarely "optimize" them. A disadvantage of optimization is that it tends to greatly complicate a program. Although significant savings may be involved, it is dangerous to apply optimization techniques at too early a stage in the development of a program. In particular, we shall not consider optimizing until we have a good assembly language implementation which we have fully analyzed, so that we can tell where the improvements will do the most good. Knuth [24] presents a fuller discussion of these issues.

Many misleading conclusions have been drawn and reported in the literature based on empirical performance statistics comparing particular implementations of particular algorithms. Empirical testing can be valuable in some situations, but, as we have seen, the structures of permutation generation algorithms are so similar that the empirical tests which have been per-
formed have really been comparisons of compilers, programmers, and computers, not of algorithms. We shall see that the differences between the best algorithms are very subtle, and they will become most apparent as we analyze the assembly language programs. Fortunately, the assembly language implementations aren't much longer than the high-level descriptions. (This turns out to be the case with many algorithms.)

## A Recursive Method (Heap)

We shall begin by looking at the implementation of Heap's method. A direct "hand compilation" of Algorithm 2 leads immediately to Program 1. The right-hand side of the program listing simply repeats the text of Algorithm 2; each statement is attached to its assembly-language equivalent.

This direct translation of Algorithm 2 is more efficient than most automatic translators would produce; it can be further improved in at least three ways. First, as we have already noted in Section 1, the test while $i \leq N$ need not be performed after we have set $i$ to 2 (if we assume $N>1$ ), so we can replace JMP WHILE in Program 1 by JMP LOOP. But this unconditional jump can be removed from the inner loop by moving the three instructions at LOOP down in its place (this is called rotating the loop-see [24]). Second, the test for whether $i$ is even or odd can be made more efficient by maintaining a separate Register X which is defined to be 1 if $i$ is even and -1 if $i$ is odd. (This improvement applies to most computers, since few have a jump if even instruction.) Third, the variable $k$ can be eliminated, and some time saved, if C(I) is updated before the ex-

Program 1. Direct Implementation of Heap's Method

| INIT | LD | Z,1 |  |
| :---: | :---: | :---: | :---: |
|  | LD | I,N | $1=N$, |
|  | ST | Z,C(I) | loop c[1] $=1$, |
|  | CMP | 1,2 |  |
|  | JLE | CALL | whlle $1>2$ |
|  | SUB | I,1 | $1=1-1$, |
|  | JMP | INIT | repeat, |
| CALL | CALL | PROCESS | process, |
| LOOP | LD | J, C(I) | loop |
|  | CMP | J, 1 |  |
|  | JE | ELSE | If $\mathrm{c}[1]<1$ |
| THEN | LD | T, I | then |
|  | AND | T,1 |  |
|  | JZ | T,EVEN | If / odd |
|  | LD | K, 1 | then $k=1$ |
|  | JMP | EXCH | else $k=c[1]$ |
| EVEN | LD | K,J | endlf, |
| EXCH | LD | T, P(l) |  |
|  | LD | T1, P(K) |  |
|  | ST | T1, P(I) |  |
|  | ST | T, P(K) | $\mathrm{P}[1]=\mathrm{P}[k]$, |
|  | ADD | J, 1 |  |
|  | ST | J,C(1) | $c[1]=c[1]+1$, |
|  | LD | 1,2 | 1 =2, |
|  | CALL | PROCESS | process, |
|  | JMP | WHILE |  |
| ELSE | ST | Z,C(I) | else $c[i] \cdot=1$, |
|  | ADD | I,1 | $t .=t+1$ endif, |
| WHILE | CMP | I,N | while $1 \leq N$ |
|  | JLE | LOOP | repeat, |

change is made, freeing Register J. These improvements lead to Program 2. (The program uses the instruction LDN $\mathrm{X}, \mathrm{X}$ which simply complements Register X.) Notice that even if we were to precompute the index table, as in Algorithm 1, we probably would not have a program as efficient as Program 2. On computers with a memory-to-memory move instruction, we might gain further efficiency by implementing the exchange in three instructions rather than in four.
Each instruction in Program 2 is labelled with the number of times it is executed when the program is run to completion. These labels are arrived at through a flow analysis which is not difficult for this program. For example, CALL PROCESS (and the two instructions preceding it) must be executed exactly $N$ ! times, since the program generates all $N$ ! permutations. The instruction at CALL can be reached via JLE CALL (which happens exactly once) or by falling through from the

PROGRAM 2. Improved Implementation of Heap's Method

| INIT | LD | Z,1 | 1 |
| :---: | :---: | :---: | :---: |
|  | LD | I,N | 1 |
|  | ST | Z, C(I) | $N-1$ |
|  | CMP | I,2 | $N-1$ |
|  | JLE | CALL | $N-1$ |
|  | SUB | I,1 | $N-1$ |
| THEN | JMP | INIT | $N-1$ |
|  | ADD | J, 1 | $N^{\prime \prime}-1$ |
|  | ST | J,C(I) | $N^{\prime}-1$ |
|  | JP | X,EXCH | N -1 |
|  | LD | J,2 | $A_{N}$ |
| EXCH | LD | T, P(I) | $\mathrm{N}^{\prime}-1$ |
|  | LD | T1, P-1(J) | $N^{\prime}-1$ |
|  | ST | T1, P( 1 ) | $N^{\prime}-1$ |
|  | ST | T, P-1(J) | $N^{\prime}-1$ |
| CALL | LD | I,2 | $N^{\prime}$ |
|  | LD | X,1 | $N^{\prime}$ |
|  | CALL | PROCESS | $N^{\prime}$ |
| LOOP | LD | J, C(l) | $N^{\prime}+B_{N}-1$ |
|  | CMP | J, | $N^{\prime}+B_{N}-1$ |
|  | JL | THEN | $N^{\prime}+B_{N}-1$ |
| ELSE | ST | Z.C(I) | $B_{N}$ |
|  | LDN | X, X | $B_{N}$ |
|  | ADD | I,1 | $B_{N}$ |
|  | CMP | I,N | $B_{N}$ |
|  | JLE | LOOP | $B_{N}$ |

preceding instruction (which therefore must happen $N$ !-1 times). Some of the instruction frequencies are more complicated (we shall analyze the quantities $A_{N}$ and $B_{N}$ in detail below), but all of the instructions can be labelled in this manner (see [31]). From these frequencies, we can calculate the total running time of the program, if we know the time taken by the individual instructions. We shall assume that instructions which reference data in memory take two time units, while jump instructions and other instructions which do not reference data in memory take one time unit. Under this model, the total running time of Program 2 is

$$
19 N!+A_{N}+10 B_{N}+6 N-20
$$

time units. These coefficients are typical, and a similar exact expression can easily be derived for any particular implementation on any particular real machine.

The improvements by which we derived Program 2 from Program 1 are applicable to most computers, but they are intended only as examples of the types of simple transformations which can lead to substantial improvements when programs are implemented. Each of the improvements results in one less instruction which is executed $N$ ! times, so its effect is significant. Such coding tricks should be applied only when an algorithm is well understood, and then only to the parts of the program which are executed most frequently. For example, the initialization loop of Program 2 could be rotated for a savings of $N-2$ time units, but we have not bothered with this because the savings is so insignificant compared with the other improvements. On the other hand, further improvements within the inner loop will be available on many computers. For example, most computers have a richer set of loop control instructions than we have used: on many machines the last three instructions in Program 2 can be implemented with a single command. In addition, we shall examine another, more advanced improvement below.

To properly determine the effectiveness of these improvements, we shall first complete the analysis of Program 2. In order to do so, we need to analyze the quantities $A_{N}$
and $B_{N}$. In the algorithm, $A_{N}$ is the number of times the test $i$ odd succeeds, and $B_{N}$ is the number of times the test $c[i]=1$ succeeds. By considering the recursive structure of the algorithm, we quickly find that the recurrence relations

$$
A_{N}=N A_{N-1}+ \begin{cases}0 & N \text { even } \\ N-1 & N \text { odd }\end{cases}
$$

and

$$
B_{N}=N B_{N-1}+1
$$

hold for $N>1$, with $A_{1}=B_{1}=0$. These recurrences are not difficult to solve. For example, dividing both sides of the equation for $B_{N}$ by $N!$ we get
$\frac{B_{N}}{N^{\prime}}=\frac{B_{N-1}}{(N-1)!}+\frac{1}{N!}=\frac{B_{N-2}}{(N-2)!}+\frac{1}{(N-1)!}+\frac{1}{N!}$

$$
=\cdots=\frac{1}{2!}+\frac{1}{3!}+\cdots+\frac{1}{N!}
$$

or

$$
B_{N}=N!\sum_{z \leq k \leq N} \frac{1}{k!} .
$$

This can be more simply expressed in terms of the base of the natural logarithms, $e$, which has the series expansion $\sum_{k \geq 0} 1 / k!$ it is easily verified that

$$
B_{N}=\lfloor N!(e-2)\rfloor
$$

That is, $B_{N}$ is the integer part of the real number $N!(e-2)$ (or $B_{N}=N!(e-2)+\epsilon$ with $0 \leq \epsilon<1$ ). The recurrences for $A_{N}$ can be solved in a similar manner to yield the result

$$
A_{N}=N!\sum_{2 s k s N} \frac{(-1)^{k}}{k!}=\{N!/ e\}
$$

Substituting these into the expression above, we find that the total running time of Program 2 is

$$
(19+(1 / e)+10(e-2)) N!+6 N+\mathrm{O}(1)
$$

or about $26.55 N$ ! time units. Table 3 shows the values of the various quantities in this analysis.

We now have a carefully implemented program whose performance we understand, and it is appropriate to consider how the program can be further "optimized." A standard technique is to identify situations that occur frequently and handle them separately in as efficient a manner as possible. For example, every other exchange performed by Program 2 is simply $\mathrm{P}[1]:=: \mathrm{P}[2]$. Rather than have the program go all the way through the main loop to discover this, incrementing and then testing $c$ [2], etc., we can gain efficiency by simply replacing $i:=2$ by $i:=3 ;$ process; $\mathrm{P}[1]:=: \mathrm{P}[2]$ in Algorithm 2. (For the purpose of this discussion assume that there is a statement $i:=2$ following the initialization loop in Algorithm 2.) In general, for any $n>1$, we can replace $i:=2$ by $\imath:=n+1$; process all permutations of $\mathrm{P}[1], \cdots, \mathrm{P}[n]$. This idea was first applied to the permutation enumeration problem by Boothroyd [2]. For small $n$, we can quite compactly write in-line code to generate all permutations of $\mathrm{P}[1], \cdots, \mathrm{P}\lfloor n\rfloor$. For example, taking $n=3$ we may simply replace

| CALL | LD | 1,2 |
| :--- | :--- | :--- |
|  | LD | $x, 1$ |
|  | CALL | PROCESS |

in Program 2 by the code in Program 3,
table 3. Analysis of Program $2\left(T_{N}=19!+A_{N}+10 B_{N}+6 N-20\right)$

| $N$ | $N!$ | $A_{N}$ | $B_{N}$ | $T_{N}$ | $26.55 N!$ |
| ---: | ---: | ---: | ---: | ---: | ---: |
|  |  |  | 0 |  |  |
| 1 | 1 | 0 | 0 | 1 | 40 |
| 2 | 6 | 2 | 4 | 154 | $56+$ |
| 3 | 24 | 8 | 17 | 638 | $159+$ |
| 4 | 120 | 44 | 86 | 3194 | $3186+$ |
| 5 | 5040 | 264 | 517 | 19130 | 19116 |
| 6 | 40320 | 1854 | 3620 | 133836 | 133812 |
| 7 | 362880 | 133396 | 28961 | 1070550 | 1070496 |
| 8 | 3628800 | 1334960 | 260650 | 9634750 | 9634464 |
| 9 | 149684570 | 2867501 | 96347210 | 96344640 |  |
| 10 | 479001600 | 176214840 | 344058145 | 1279818936 | 1059791040 |
| 11 |  |  | 1271726742 | 12717492480 |  |

which efficiently permutes $\mathrm{P}[1], \mathrm{P}[2], \mathrm{P}[3]$. (While only the code that differs from Program 2 is given here, "Program 3 " refers to the entire improved program.)

The analysis of Program 3 differs only slightly from that of Program 2. This is fortunate, for it is often difficult to determine the exact effect of such major improvements. First, each of the new instructions is clearly executed $N!/ 6$ times, and each occurrence of $N$ ! in Program 2's frequencies becomes $N!/ 6$ for Program 3; thus, the total running time is

$$
(50 / 6) N!+A^{\prime}{ }_{N}+B_{N}^{\prime}+6 N-20 .
$$

Next, the analysis for $A_{N}$ and $B_{N}$ given above still holds, except that the initial conditions are different. We find that

$$
\begin{aligned}
& \left.A_{N}^{\prime}=N!\sum_{4 \leqslant \leqslant \leqslant N} \frac{(-1)^{k}}{k!}=\left[N!\left(\frac{1}{e}-\frac{1}{3}\right)\right] \right\rvert\, \\
& B_{N}^{\prime}=N!\sum_{4 \leqslant \leqslant \leqslant N} \frac{1}{k!}=\left[N!\left(e-\frac{8}{3}\right)\right]
\end{aligned}
$$

and the total running time of Program 3 is then about 8.88 N !.
By taking larger values of $n$ we can get further improvements, but at the cost of

PROGRAM 3. Optimized Inner Loop for Program 2

| CALL | LD | 1,4 |
| :--- | :--- | :--- |
|  | LD | X, 1 |
|  | CALL | PROCESS |
|  | LD | T1,P(1) |
|  | LD | T2,P(2) |
|  | LD | T3,P(3) |
|  | ST | T1,P(2) |
|  | ST | T2,P(1) |
|  | CALL | PROCESS |
|  | ST | T3,P(1) |
|  | ST | T2,P(3) |
|  | CALL | PROCESS |
|  | ST | T1,P(1) |
|  | ST | T3,P(2) |
|  | CALL | PROCESS |
|  | ST | T2,P(1) |
|  | ST | T1,P(3) |
|  | CALL | PROCESS |
|  | ST | T3,P(1) |
|  | ST | T2,P(2) |
| CALL | PROCESS |  |

$n+3 n!$ lines of code. This is an example of a space-time tradeoff where the time saved is substantial when $n$ is small, but the space consumed becomes substantial when $n$ is large. For $n=4$, the total running time goes down to about $5.88 N$ ! and it is probably not worthwhile to go further, since the best that we could hope for would be 5 N ! (the cost of two stores and a call).

On most computers, if Program 2 is "optimized" in the manner of Program 3 with $n=4$, Heap's method will run faster than any other known method.

## An Iterative Method (Ives)

The structures of Algorithm 2 and Algorithm 4 are very similar, so that a direct "hand compilation" of Ives' method looks very much like Program 1. By rotating the loop and maintaining the value $N+1-i$ in a separate register we get Program 4, an improved implementation of Ives' method which corresponds to the improved implementation of Heap's method in Program 2.
The total running time of this program is

$$
18 N!+21 D_{N}+10 N-25,
$$

where $D_{N}$ is the number of times $i:=i+1$ is executed in Algorithm 4. Another quantity, $C_{N}$, the number of times the test $\mathrm{P}[N+1-i]=\mathrm{Q}[N+1-i]$ fails, happens to cancel out when the total running time is computed. These quantities can be analyzed in much the same way that $A_{N}$ and $B_{N}$ were analyzed for Program 2: they satisfy the recurrences

$$
\begin{aligned}
& C_{N}=C_{N-2}+(N-1)!-(N-2)! \\
& D_{N}=D_{N-2}+(N-2)!
\end{aligned}
$$

so that
$C_{N}=(N-1)!-(N-2)!+(N-3)!-(N-4)!+\cdots$
$D_{N}=(N-2)!+(N-4)!+(N-6)!+\cdots$
and the total running time of Program 2 is

$$
18 N!+21(N-2)+O((N-4)!)
$$

or about

$$
\left(18^{\circ}+\frac{21}{N(N-1)}\right) N!
$$

Thus Program 4 is faster than Program 2: the improved implementation of Ives' method uses less overhead per permuta-

PROGRAM 4 Improved Implementation of Ives'

tion than the improved implementation of Heap's method, mainly because it does less counter manipulation. Other iterative methods, like the Johnson-Trotter algorithm (or the version of Ives' method, Algorithm 4a, which does not require the elements to be distinct), are only slightly faster than Heap's method.

However, the iterative methods cannot be optimized quite as completely as we were able to improve Heap's method. In Algorithm 4 and Program 4, the most frequent operation is $\mathrm{P}[c[N]]:=: \mathrm{P}[c[N]+1]$; $c[N]:=c[N]+1$; all but $1 / N$ of the exchanges are of this type. Therefore, we should program this operation separately. (This idea was used by Ehrlich [10, 11].) Program 4 can be improved by inserting the code given in Program 5 directly after

[^4]PROCESS
(As before, we shall write down only the new code, but make reference to the entire optimized program as "Program 5".) In this program, Pointer J is kept negative so that we can test it against zero, which can be done efficiently on many computers. Alternatively, we could sweep in the other direction, and have J range from $N-1$ to 0 . Neither of these tricks may be necessary on computers with advanced loop control instructions.

To find the total running time of Program 5, it turns out that we need only replace $N$ ! by ( $N-2$ )! everywhere in the frequencies in Program 4, and then add the frequencies of the new instructions. The result is

$$
9 N!+2(N-1)!+18(N-2)!+\mathrm{O}((N-4)!),
$$

not quite as fast as the "optimized" version of Heap's algorithm (Program 3). For a fixed value of $N$, we could improve the program further by completely unrolling the inner loop of Program 5. The second through eighth instructions of Program 5 could be replaced by

| LD | T,P+1 |
| :--- | :--- |
| ST | T,P |
| ST | $V, P+1$ |
| CALL | PROCESS |
| LD | T,P+2 |
| ST | $\mathrm{T}, \mathrm{P}+1$ |
| ST | $\mathrm{V}, \mathrm{P}+2$ |
| CALL | PROCESS |
| LD | T,P+3 |
| ST | T,P+2 |
| ST | $V, P+3$ |
| CALL | PROCESS |

(This could be done, for example, by a macro generator). This reduces the total running time to

$$
7 N!+(N-1)!+18(N-2)!+\mathrm{O}((N-4)!)
$$

which is not as fast as the comparable highly optimized version of Heap's method (with $n=4$ ).

It is interesting to note that the optimization technique which is appropriate for the recursive programs (handling small cases separately) is much more effective than the optimization technique which is

| PROGRAM 5 |  | Optimized Inner Loop for Program 4 |  |
| :---: | :---: | :---: | :---: |
| EXLP | : |  |  |
|  | CALL | PROCESS | $(N-1)^{\prime}$ |
|  | LD | J,1-N | $(\mathrm{N}-1)^{\prime}$ |
| INLP | LD | T, $\mathbf{P}+\mathbf{N + 1}(\mathrm{J})$ | $N^{\prime}-(N-1)^{\prime}$ |
|  | ST | $\mathrm{T}, \mathrm{P}+\mathrm{N}(\mathrm{J})$ | $N^{\prime}-(N-1)^{\prime}$ |
|  | ST | $\mathrm{V}, \mathrm{P}+\mathrm{N}+1(\mathrm{~J})$ | $N^{\prime}-(N-1)^{\prime}$ |
|  | CALL | PROCESS | $N^{\prime}-(N-1)^{\prime}$ |
|  | ADD | J, 1 | $\mathrm{N}^{\prime}-(N-1)^{\prime}$ |
|  | JN | J,INLP | $N^{\prime}-(N-1)^{\prime}$ |
|  | LD | T, P+1 | $(N-1)^{\prime}$ |
|  | ST | T, P+N | $(N-1)^{\prime}$ |
|  | ST | $V, P+1$ | $(N-1)^{\prime}$ |
|  | CMP | T, $\mathrm{Q}+\mathrm{N}$ | $(N-1)^{\prime}$ |
|  | JNE | EXLP | $(\mathrm{N}-1)^{\prime}$ |

appropriate for the iterative programs (loop unrolling).

## A Cyclic Method (Langdon)

It is interesting to study Langdon's cyclic method (Algorithm 6) in more detail, because it can be implemented with only a few instructions on many computers. In addition, it can be made to run very fast on computers with hardware rotation capabilities.
To implement Algorithm 6, we shall use a new instruction

```
MOVE TO, FROM(I)
```

which, if Register I contains the number $i$, moves $l$ words starting at Location FROM to Location TO. That is, the above instruction is equivalent to

|  | LD J,0 |
| :--- | :--- |
| LOOP | T,FROM(J) |
|  | T,TO(J) |
|  | ADD J,1 |
|  | CMPJ,I |
|  | JL LOOP |

We shall assume that memory references are overlapped, so that the instruction takes $2 i$ time units. Many computers have "block transfer" instructions similar to this, although the details of implementation vary widely.
For simplicity, let us further suppose that $\mathrm{P}[1], \cdots, \mathrm{P}[N]$ are initially the integers $0,1, \cdots, N-1$, so that we don't have to bother with the Q array of Algorithm 6.

With these assumptions, Langdon's method is particularly simple to implement, as shown in Program 6. Only eight assembly language instructions will suffice on many computers to generate all permutations of $\{0,1, \cdots, N-1\}$.

As we have already noted, however, the moves tend to be long, and the method is not particularly efficient if $N$ is not small. Proceeding as we have before, we see that $E_{N}$ and $F_{N}$ satisfy

$$
\begin{aligned}
& \boldsymbol{E}_{N}=\sum_{1, \leqslant k N N_{-1}} k! \\
& \boldsymbol{F}_{N}=\sum_{1 \leq k \leqslant s, v} k k!=(N+1)!-1
\end{aligned}
$$

(Here $F_{N}$ is not the frequency of execution of the MOVE instruction, but the total number of words moved by it.) The total running time of Program 6 turns out to be

$$
N!\left(2 N+10+\frac{9}{N}\right)+\left(0(N-2)^{\prime}\right)
$$

It is faster than Program 2 for $N<8$ and faster than Program 4 for $N<4$, but it is much slower for larger $N$.

By almost any measure, Program 6 is the simplest of the programs and algorithms that we have seen so far. Furthermore, on most computer systems it will run faster than any of the algorithms implemented in a high-level language. The algorithm fueled a controversy of sorts (see other references in [25]) when it was first introduced, based on just this issue.
Furthermore, if hardware rotation is available, Program 6 may be the method of choice. Since ( $N-1$ )/ $N$ of the rotations are of length $N$, the program may be optimized in the manner of Program 5 around a fourinstruction inner loop (call, rotate, compare, conditional jump). On some ma-

PROGRAM 6 Implementation of Langdon's Method

| Method |  |  |  |
| :--- | :--- | :--- | :--- |
| THEN | LD | I,N-1 | $N^{\prime}$ |
|  | CALL | PROCESS | $N^{\prime}$ |
| LOOP | LD | T,P+1 | $N^{\prime}+E_{N}$ |
|  | MOVE | P,P+1(I) | $F_{N}$ |
|  | ST | T,P+1(I) | $N^{\prime}+E_{N}$ |
|  | CMP | T,I | $N^{\prime}+E_{N}$ |
|  | JNE | THEN | $N^{\prime}+E_{N}$ |
|  | SUB | 1,1 | $E_{N}$ |
|  | JNZ | LOOP | $E_{N}$ |

chines, the rotate might be performed in, say, two time units (for example, if parallelism were available, or if $P$ were maintained in registers), which would lead to a total time of $5 N!+\mathrm{O}((N-1)!)$. We have only sketched details here because the issues are so machine-dependent: the obvious point is that exotic hardware features can have drastic effects upon the choice of algorithm.

## CONCLUSION

The remarkable similarity of the many permutation enumeration algorithms which have been published has made it possible for us to draw some very definite conclusions regarding their performance. In Section 1, we saw that the method given by Heap is slightly simpler (and therefore slightly more efficient) than the methods of Wells and Boothroyd, and that the method given by Ives is simpler and more efficient than the methods of Johnson and Trotter (and Ehrlich). In Section 2, we found that the cyclic and lexicographic algorithms will not compete with these, except possibly for Langdon's method, which avoids some of the overhead in the control structure inherent in the methods. By carefully implementing these algorithms in Section 3 and applying standard code optimization techniques, we found that Heap's method will run fastest on most computers, since it can be coded so that most permutations are generated with only two store instructions.
However, as discussed in the Introduction, our accomplishments must be kept in perspective. An assembly-language implementation such as Program 3 may run 50 to 100 times faster than the best previously published algorithms (in high-level languages) on most computer systems, but this means merely that we can now generate all permutations of 12 elements in one hour of computer time, where before we could not get to $N=11$. On the other hand, if we happen to be interested only in all permutations of 10 elements, we can now get them in only 15 seconds, rather than 15 minutes.
The problem of comparing different algorithms for the same task arises again
and again in computer science, because new algorithms (and new methods of expressing algorithms) are constantly being developed. Normally, the kind of detailed analysis and careful implementation done in this paper is reserved for the most important algorithms. But permutation generation nicely illustrates the important issues. An appropriate choice between algorithms for other problems can be made by studying their structure, implementation, analysis, and "optimization" as we have done for permutation generation.

## ACKNOWLEDGMENTS

Thanks are due to $P$ Flanagan, who implemented and checked many of the algorithms and programs on a real computer. Also, I must thank the many authors listed below for providing me with such a wealth of material, and I must apologize to those whose work I may have misunderstood or misrepresented. Finally, the editor, P. J. Denning, must be thanked for his many comments and suggestions for improving the readability of the manuscript.

## REFERENCES

[1] Boothroyd, J "PERM (Algorithm 6)," Computer Bulletin 9, 3 (Dec. 1965), 104
[2] Boothroyd, J. "Permutation of the elements of a vector (Algorithm 29)"; and "Fast permutation of the elements of a vector (Algorithm 30)," Computer I. 10 (1967), 310-312.
[3] Bratley, P "Permutations with repetitions (Algorithm 306)," Comm ACM 10, 7 (July 1967), 450
[4] Coveyou, R R.; and Sullivan, J. G. "Permutation (Algorithm 71)," Comm. ACM 4, 11 (Nov 1961), 497.
[5] Dershowitz, N. "A simplified loop-free algorithm for generating permutations," BIT 15 (1975), 158-164.
[6] Dijkstra, E W A discipline of programming, Prentice-Hall, Englewood Cliffs, N J , 1976.
[7] Dijkstra, E. W "On a gauntlet thrown by David Gries," Acta Informatica 6, 4 (1976), 357.
[8] Durstenfeld, R. "Random permutation (Algorithm 235)," Comm. ACM 7, 7 (July 1964), 420.
[9] Eaves, B. C "Permute (Algorithm 130)," Comm. ACM 5, 11 (Nov. 1962), 551 (See also: remarks by R. J. Ord-Smith in Comm ACM 10, 7 (July, 1967), 452-3)
[10] Ehrlich, G. "Loopless algorithms for generating permutations, combinations and other combinatorial configurations," $I$ ACM 20, 3 (July 1973), 500-513.
[11] Ehrlich, G "Four combinatorial algorithms (Algorithm 466)," Comm ACM 16, 11 (Nov 1973), 690-691.
[12] Even, S. Algorithmic combinatorics, Macmullan, Inc., N.Y., 1973.
[13] Fike, C T. "A permutation generation method," Computer f 18, 1 (Feb. 1975), 21-22.
[14] Fischer, L. L.; and Krause, K. C , Lehrbuch der Combinatıonslehre und der Artthmetik, Dresden, 1812
[15] Hall, M.; and Knuth, D. E. "Combinatorial analysis and computers," American Math. Monthly 72, 2 (Feb. 1965, Part II), 21-28.
[16] Heap, B R. "Permutations by interchanges," Computer J. 6 (1963), 293-4.
[17] Howell, J R. "Generation of permutations by addtion," Math. Comp 16 (1962), 243-44
[18] Howell, J. R. "Permutation generator (Algorithm 87)," Comm. ACM 5, 4 (April 1962), 209. (See also: remarks by R. J. Ord-Smith in Comm ACM 10, 7 (July 1967), 452-3.)
[19] Ives, F. M. "Permutation enumeration: four new permutation algorithms," Comm. ACM 19, 2 (Feb. 1976), 68-72.
[20] Johnson, S. M. "Generation of permutations by adjacent transposition," Math. Comp. 17 (1963), 282-285.
[21] Knuth, D. E. "Fundamental algorithms," in The art of computer programming 1, AddisonWesley, Co., Inc., Reading, Mass., 1968.
[22] Knuth, D. E. "Seminumerical algorithms," in The art of computer programming 2, Addi-son-Wesley, Co., Inc., Reading, Mass., 1969.
[23] Knuth, D. E "Sorting and searching," in The art of computer programming 3, AddisonWesley, Co., Inc, Reading, Mass, 1972
[24] Knuth, D. E. "Structured programming with go to statements," Computıng Surveys 6, 4 (Dec. 1974), 261-301.
[25] Langdon, G. G., Jr., "An algorithm for generating permutations," Comm. ACM 10, 5 (May 1967), 298-9. (See also. letters by R. J. Ord-Smith in Comm. ACM 10, 11 (Nov. 1967), 684; by B. E. Rodden in Comm. ACM 11, 3 (March 1968), 150; CR Rev. 13,891, Computing Reviews 9, 3 (March 1968), and letter by Langdon in Comm. ACM 11, 6 (June 1968), 392.)
[26] Lehmer, D H. "Teaching combinatorial tricks to a computer," in Proc. of Symposum Appl. Math Combinatorial Analysis, Yol. 10, American Mathematical Society, Providence, R.I , 1960, 179-193.
[27] Lehmer, D. H. "The machine tools of combinatorics," in Applied combinatorzal mathematccs (E. F. Beckenbach, [Ed.]), John Wiley, \& Sons, Inc., N Y , 1964
[28] Lehmer, D. H. "Permutation by adjacent interchanges," Amercan Math. Monthly 72, 2 (Feb. 1965, Part II), 36-46.
[29] Ord-Smith, R J. "Generation of permutation sequences Part 1," Computer J. 13, 3 (March 1970), 152-155.
[30] Ord-Smith, R. J. "Generation of permutation sequences: Part 2," Computer J. 14, 2 (May 1971), 136-139.
[31] Ord-Smith, R. J. "Generation of permutations in pseudo-lexicographic order (Algorithm 308)," Comm ACM 10, 7 (July 1967),

452 (See also: remarks in Comm ACM 12, 11 (Nov 1969), 638.)
[32] Ord-Smith, R J "Generation of permutations in lexicographic order (Algorithm 323)," Comm. ACM 11, 2 (Feb. 1968), 117 (See also: certification by I. M. Leitch in Comm ACM 12, 9 (Sept. 1969), 512.)
[33] Peck, J E. L; and Schrack, G F "Permute (Algorithm 86)," Comm. ACM 5, 4 (April 1962), 208.
[34] Phillips, J. P. N. "Permutation of the elements of a vector in lexicographic order (Algorithm 28)," Computer J. 10 (1967), 310-311
[35] Pleszczynski, S. "On the generation of permutations," Information Processing Letters 3, 6 (July 1975), 180-183.
[36] Riordan, J. An introduction to combinatorual analysts, John Wiley \& Sons, Inc., N.Y., 1958.
[37] Rohl, J. S., "Programming improvements to Fike's algorithm for generating permutations," Computer J. 19, 2 (May 1976), 156
[38] Robinson, C. L. "Permutation (Algorithm 317)," Comm. ACM 10, 11 (Nov. 1967), 729.
[39] SAG, T. W. "Permutations of a set with repetitions (Algorithm 242)," Comm ACM 7, 10 (Oct 1964), 585.
[40] Schrak, G. F; and Shimrat, M "Permutation in lexicographic order (Algorithm 102)," Comm. ACM 5, 6 (June 1962), 346. (See also: remarks by $R$ J Ord-Smith in Comm ACM 10, 7 (July 1967), 452-3.)
[41] Shen, M.-K "On the generation of permutations and combinatıons," BIT 2 (1962), 228231
[42] Shen, M -K "Generation of permutations in lexicographic order (Algorithm 202)," Comm ACM 6, 9 (Sept. 1963), 517. (See also: remarks by R. J. Ord-Smith in Comm. ACM 10, 7 (July 1967), 452-3.)
[43] Sloane, N. J. A. A handbook of integer sequences, Academic Press, Inc , N Y , 1973
[44] Tompkins, C. "Machine attacks on problems whose variables are permutations," in Proc Symposium in Appl. Math., Numerical Analysus, Vol 6, McGraw-Hıll, Inc., N.Y., 1956, 195-211
[45] Trotter, H. F. "Perm (Algorithm 115)," Comm. ACM 5, 8 (August 1962), 434-435.
[46] Walker, R. J "An enumerative technique for a class of combinatorial problems," in Proc Symposium in Appl. Math, Combinatorial Analysis, Vol. 10, American Mathematical Soclety, Providence, R.I , 1960, 91
[47] Wells, M. B "Generation of permutations by transposition," Math Comp 15 (1961), 192195.
[48] Wells, M. B. Elements of combinatorial computing, Pergamon Press, Elmsford, N.Y., 1971.
[49] Whitehead, E. G. Combinatorial algorithms. (Notes by C. Frankfeldt and A. E. Kaplan) Courant Institute, New York Univ., 1973, 11-17.


[^0]:    *This work was supported by the Natıonal Science Foundation Grant No. MCS75-23738

[^1]:    Copyright © 1977, Association for Computing Machinery, Inc. General permission to republish, but not for profit, all or part of this material is granted provided that ACM's copyright notice is given and that reference is made to the publication, to its date of issue, and to the fact that reprinting privileges were granted by permission of the Association for Computing Machinery

[^2]:    Algorithm 5 (Tompkins-Paige)
    $i .=N$; loop: $c[i]=1$ while $t>2: \iota:=u-1$ repeat; process;

[^3]:    $\imath:=1$; loop: $\mathrm{Q}[z]:=\mathrm{P} \mid l]$ while $l<N . \imath:=\imath+1$ repeat, process;

[^4]:    CALL

