

A Fast String Searching Algorithm

Robert S. Boyer
Stanford Research Institute

J Strother Moore
Xerox Palo Alto Research Center

An algorithm is presented that searches for the location, "*i*," of the first occurrence of a character string, "*pat*," in another string, "*string*." During the search operation, the characters of *pat* are matched starting with the last character of *pat*. The information gained by starting the match at the end of the pattern often allows the algorithm to proceed in large jumps through the text being searched. Thus the algorithm has the unusual property that, in most cases, not all of the first *i* characters of *string* are inspected. The number of characters actually inspected (on the average) decreases as a function of the length of *pat*. For a random English pattern of length 5, the algorithm will typically inspect $i/4$ characters of *string* before finding a match at *i*. Furthermore, the algorithm has been implemented so that (on the average) fewer than $i + patlen$ machine instructions are executed. These conclusions are supported with empirical evidence and a theoretical analysis of the average behavior of the algorithm. The worst case behavior of the algorithm is linear in $i + patlen$, assuming the availability of array space for tables linear in *patlen* plus the size of the alphabet.

Key Words and Phrases: bibliographic search, computational complexity, information retrieval, linear time bound, pattern matching, text editing

CR Categories: 3.74, 4.40, 5.25

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.

Authors' present addresses: R.S. Boyer, Computer Science Laboratory, Stanford Research Institute, Menlo Park, CA 94025. This work was partially supported by ONR Contract N00014-75-C-0816; JS. Moore was in the Computer Science Laboratory, Xerox Palo Alto Research Center, Palo Alto, CA 94304, when this work was done. His current address is Computer Science Laboratory, SRI International, Menlo Park, CA 94025.

1. Introduction

Suppose that *pat* is a string of length *patlen* and we wish to find the position *i* of the leftmost character in the first occurrence of *pat* in some string *string*:

```
pat:          AT-THAT
string: ... WHICH-FINALLY-HALTS.--AT-THAT-POINT ...
i:           ↑
```

The obvious search algorithm considers each character position of *string* and determines whether the successive *patlen* characters of *string* starting at that position match the successive *patlen* characters of *pat*. Knuth, Morris, and Pratt [4] have observed that this algorithm is quadratic. That is, in the worst case, the number of comparisons is on the order of $i * patlen$.¹

Knuth, Morris, and Pratt have described a linear search algorithm which preprocesses *pat* in time linear in *patlen* and then searches *string* in time linear in $i + patlen$. In particular, their algorithm inspects each of the first $i + patlen - 1$ characters of *string* precisely once.

We now present a search algorithm which is usually "sublinear": It may not inspect each of the first $i + patlen - 1$ characters of *string*. By "usually sublinear" we mean that the expected value of the number of inspected characters in *string* is $c * (i + patlen)$, where $c < 1$ and gets smaller as *patlen* increases. There are patterns and strings for which worse behavior is exhibited. However, Knuth, in [5], has shown that the algorithm is linear even in the worst case.

The actual number of characters inspected depends on statistical properties of the characters in *pat* and *string*. However, since the number of characters inspected on the average decreases as *patlen* increases, our algorithm actually speeds up on longer patterns.

Furthermore, the algorithm is sublinear in another sense: It has been implemented so that on the average it requires the execution of fewer than $i + patlen$ machine instructions per search.

The organization of this paper is as follows: In the next two sections we give an informal description of the algorithm and show an example of how it works. We then define the algorithm precisely and discuss its efficient implementation. After this discussion we present the results of a thorough test of a particular machine code implementation of our algorithm. We compare these results to similar results for the Knuth, Morris, and Pratt algorithm and the simple search algorithm. Following this empirical evidence is a theoretical analysis which accurately predicts the performance measured. Next we describe some situations in which it may not be advantageous to use our algorithm. We conclude with a discussion of the history of our algorithm.

¹ The quadratic nature of this algorithm appears when initial substrings of *pat* occur often in *string*. Because this is a relatively rare phenomenon in string searches over English text, this simple algorithm is *practically* linear in $i + patlen$ and therefore acceptable for most applications.

2. Informal Description

The basic idea behind the algorithm is that more information is gained by matching the pattern from the right than from the left. Imagine that *pat* is placed on top of the left-hand end of *string* so that the first characters of the two strings are aligned. Consider what we learn if we fetch the *patlen*th character, *char*, of *string*. This is the character which is aligned with the last character of *pat*.

Observation 1. If *char* is known not to occur in *pat*, then we know we need not consider the possibility of an occurrence of *pat* starting at *string* positions 1, 2, . . . or *patlen*: Such an occurrence would require that *char* be a character of *pat*.

Observation 2. More generally, if the last (right-most) occurrence of *char* in *pat* is δ_1 characters from the right end of *pat*, then we know we can slide *pat* down δ_1 positions without checking for matches. The reason is that if we were to move *pat* by less than δ_1 , the occurrence of *char* in *string* would be aligned with some character it could not possibly match: Such a match would require an occurrence of *char* in *pat* to the right of the rightmost.

Therefore unless *char* matches the last character of *pat* we can move past δ_1 characters of *string* without looking at the characters skipped; δ_1 is a function of the character *char* obtained from *string*. If *char* does not occur in *pat*, δ_1 is *patlen*. If *char* does occur in *pat*, δ_1 is the difference between *patlen* and the position of the rightmost occurrence of *char* in *pat*.

Now suppose that *char* matches the last character of *pat*. Then we must determine whether the previous character in *string* matches the second from the last character in *pat*. If so, we continue backing up until we have matched all of *pat* (and thus have succeeded in finding a match), or else we come to a mismatch at some new *char* after matching the last *m* characters of *pat*.

In this latter case, we wish to shift *pat* down to consider the next plausible juxtaposition. Of course, we would like to shift it as far down as possible.

Observation 3(a). We can use the same reasoning described above—based on the mismatched character *char* and δ_1 —to slide *pat* down *k* so as to align the two known occurrences of *char*. Then we will want to inspect the character of *string* aligned with the last character of *pat*. Thus we will actually shift our attention down *string* by $k + m$. The distance *k* we should slide *pat* depends on where *char* occurs in *pat*. If the rightmost occurrence of *char* in *pat* is to the right of the mismatched character (i.e. within that part of *pat* we have already passed) we would have to move *pat* backwards to align the two known occurrences of *char*. We would not want to do this. In this case we say that δ_1 is “worthless” and slide *pat* forward by $k = 1$ (which is always sound). This shifts our attention down *string* by $1 + m$. If the rightmost occurrence of *char* in

pat is to the left of the mismatch, we can slide forward by $k = \delta_1(\text{char}) - m$ to align the two occurrences of *char*. This shifts our attention down *string* by $\delta_1(\text{char}) - m + m = \delta_1(\text{char})$.

However, it is possible that we can do better than this.

Observation 3(b). We know that the next *m* characters of *string* match the final *m* characters of *pat*. Let this substring of *pat* be *subpat*. We also know that this occurrence of *subpat* in *string* is preceded by a character (*char*) which is different from the character preceding the terminal occurrence of *subpat* in *pat*. Roughly speaking, we can generalize the kind of reasoning used above and slide *pat* down by some amount so that the discovered occurrence of *subpat* in *string* is aligned with the rightmost occurrence of *subpat* in *pat* which is not preceded by the character preceding its terminal occurrence in *pat*. We call such a reoccurrence of *subpat* in *pat* a “plausible reoccurrence.” The reason we said “roughly speaking” above is that we must allow for the rightmost plausible reoccurrence of *subpat* to “fall off” the left end of *pat*. This is made precise later.

Therefore, according to Observation 3(b), if we have matched the last *m* characters of *pat* before finding a mismatch, we can move *pat* down by *k* characters, where *k* is based on the position in *pat* of the rightmost plausible reoccurrence of the terminal substring of *pat* having *m* characters. After sliding down by *k*, we want to inspect the character of *string* aligned with the last character of *pat*. Thus we actually shift our attention down *string* by $k + m$ characters. We call this distance δ_2 , and we define δ_2 as a function of the position *j* in *pat* at which the mismatch occurred. *k* is just the distance between the terminal occurrence of *subpat* and its rightmost plausible reoccurrence and is always greater than or equal to 1. *m* is just *patlen* - *j*.

In the case where we have matched the final *m* characters of *pat* before failing, we clearly wish to shift our attention down *string* by $1 + m$ or $\delta_1(\text{char})$ or $\delta_2(j)$, according to whichever allows the largest shift. From the definition of δ_2 as $k + m$ where *k* is always greater than or equal to 1, it is clear that δ_2 is at least as large as $1 + m$. Therefore we can shift our attention down *string* by the maximum of just the two δ s. This rule also applies when $m = 0$ (i.e. when we have not yet matched any characters of *pat*), because in that case $j = \text{patlen}$ and $\delta_2(j) \geq 1$.

3. Example

In the following example we use an “↑” under *string* to indicate the current *char*. When this “pointer” is pushed to the right, imagine that it drags the right end of *pat* with it (i.e. imagine *pat* has a hook on its right end). When the pointer is moved to the left, keep *pat* fixed with respect to *string*.

```

pat:           AT-THAT
string: ... WHICH-FINALLY-HALTS.--AT-THAT-POINT ...
                ↑

```

Since “F” is known not to occur in *pat*, we can appeal to Observation 1 and move the pointer (and thus *pat*) down by 7:

```

pat:           AT-THAT
string: ... WHICH-FINALLY-HALTS.--AT-THAT-POINT ...
                ↑

```

Appealing to Observation 2, we can move the pointer down 4 to align the two hyphens:

```

pat:           AT-THAT
string: ... WHICH-FINALLY-HALTS.--AT-THAT-POINT ...
                ↑

```

Now *char* matches its opposite in *pat*. Therefore we step left by one:

```

pat:           AT-THAT
string: ... WHICH-FINALLY-HALTS.--AT THAT POINT ...
                ↑

```

Appealing to Observation 3(a), we can move the pointer to the right by 7 positions because “L” does not occur in *pat*.² Note that this only moves *pat* to the right by 6.

```

pat:           AT-THAT
string: ... WHICH-FINALLY-HALTS.--AT-THAT-POINT ...
                ↑

```

Again *char* matches the last character of *pat*. Stepping to the left we see that the previous character in *string* also matches its opposite in *pat*. Stepping to the left a second time produces:

```

pat:           AT-THAT
string: ... WHICH-FINALLY-HALTS.--AT-THAT-POINT ...
                ↑

```

Noting that we have a mismatch, we appeal to Observation 3(b). The δ_2 move is best since it allows us to push the pointer to the right by 7 so as to align the discovered substring “AT” with the beginning of *pat*.³

```

pat:           AT-THAT
string: ... WHICH-FINALLY-HALTS.--AT-THAT-POINT ...
                ↑

```

This time we discover that each character of *pat* matches the corresponding character in *string* so we have found the pattern. Note that we made only 14 references to *string*. Seven of these were required to confirm the final match. The other seven allowed us to move past the first 22 characters of *string*.

² Note that δ_2 would allow us to move the pointer to the right only 4 positions in order to align the discovered substring “T” in *string* with its second from last occurrence at the beginning of the word “THAT” in *pat*.

³ The δ_1 move only allows the pointer to be pushed to the right by 4 to align the hyphens.

4. The Algorithm

We now specify the algorithm. The notation $pat(j)$ refers to the *j*th character in *pat* (counting from 1 on the left).

We assume the existence of two tables, δ_1 and δ_2 . The first has as many entries as there are characters in the alphabet. The entry for some character *char* will be denoted by $\delta_1(char)$. The second table has as many entries as there are character positions in the pattern. The *j*th entry will be denoted by $\delta_2(j)$. Both tables contain non-negative integers.

The tables are initialized by preprocessing *pat*, and their entries correspond to the values δ_1 and δ_2 referred to earlier. We will specify their precise contents after it is clear how they are to be used.

Our search algorithm may be specified as follows:

```

stringlen ← length of string.
i ← patlen.
top:  if i > stringlen then return false.
      j ← patlen.
loop: if j = 0 then return j + 1.
      if string(i) = pat(j)
      then
        j ← j - 1.
        i ← i - 1.
        goto loop.
      close;
      i ← i + max(δ1(string(i)), δ2(j)).
      goto top.

```

If the above algorithm returns false, then *pat* does not occur in *string*. If the algorithm returns a number, then it is the position of the left end of the first occurrence of *pat* in *string*.

The δ_1 table has an entry for each character *char* in the alphabet. The definition of δ_1 is:

$\delta_1(char)$ = If *char* does not occur in *pat*, then *patlen*; else *patlen* - *j*, where *j* is the maximum integer such that $pat(j) = char$.

The δ_2 table has one entry for each of the integers from 1 to *patlen*. Roughly speaking, $\delta_2(j)$ is (a) the distance we can slide *pat* down so as to align the discovered occurrence (in *string*) of the last *patlen* - *j* characters of *pat* with its rightmost plausible reoccurrence, plus (b) the additional distance we must slide the “pointer” down so as to restart the process at the right end of *pat*. To define δ_2 precisely we must define the rightmost plausible reoccurrence of a terminal substring of *pat*. To this end let us make the following conventions: Let \$ be a character that does not occur in *pat* and let us say that if *i* is less than 1 then $pat(i)$ is \$. Let us also say that two sequences of characters $[c_1 \dots c_n]$ and $[d_1 \dots d_n]$ “unify” if for all *i* from 1 to *n* either $c_i = d_i$ or $c_i = \$$ or $d_i = \$$.

Finally, we define the position of the rightmost plausible reoccurrence of the terminal substring which starts at position *j* + 1, $rpr(j)$, for *j* from 1 to *patlen*, to be the greatest *k* less than or equal to *patlen* such that

$[pat(j + 1) \dots pat(patlen)]$ and $[pat(k) \dots pat(k + patlen - j - 1)]$ unify and either $k \leq 1$ or $pat(k - 1) \neq pat(j)$.⁴ (That is, the position of the rightmost plausible reoccurrence of the substring *subpat*, which starts at $j + 1$, is the rightmost place where *subpat* occurs in *pat* and is not preceded by the character $pat(j)$ which precedes its terminal occurrence—with suitable allowances for either the reoccurrence or the preceding character to fall beyond the left end of *pat*. Note that $rpr(j)$ may be negative because of these allowances.)

Thus the distance we must slide *pat* to align the discovered substring which starts at $j + 1$ with its rightmost plausible reoccurrence is $j + 1 - rpr(j)$. The distance we must move to get back to the end of *pat* is just $patlen - j$. $delta_2(j)$ is just the sum of these two. Thus we define $delta_2$ as follows:

$$delta_2(j) = patlen + 1 - rpr(j).$$

To make this definition clear, consider the following two examples:

<i>j</i> :	1	2	3	4	5	6	7	8	9
<i>pat</i> :	A	B	C	X	X	X	A	B	C
$delta_2(j)$:	14	13	12	11	10	9	11	10	1

<i>j</i> :	1	2	3	4	5	6	7	8	9
<i>pat</i> :	A	B	Y	X	C	D	E	Y	X
$delta_2(j)$:	17	16	15	14	13	12	7	10	1

5. Implementation Considerations

The most frequently executed part of the algorithm is the code that embodies Observations 1 and 2. The following version of our algorithm is equivalent to the original version provided that $delta_0$ is a table containing the same entries as $delta_1$ except that $delta_0(pat(patlen))$ is set to an integer *large* which is greater than $stringlen + patlen$ (while $delta_1(pat(patlen))$ is always 0).

```

stringlen ← length of string.
i ← patlen.
if i > stringlen then return false.
fast: i ← i + delta_0(string(i)).
      if i ≤ stringlen then goto fast.
undo: if i ≤ large then return false.
      i ← (i - large) - 1.
      j ← patlen - 1.
slow: if j = 0 then return i + 1.
      if string(i) = pat(j)
         then
            j ← j - 1.
            i ← i - 1.
            goto slow.
         close;
      i ← i + max(delta_1(string(i)), delta_2(j)).
      goto fast.

```

⁴ Note that when $j = patlen$, the two sequences $[pat(patlen + 1) \dots pat(patlen)]$ and $[pat(k) \dots pat(k - 1)]$ are empty and therefore unify. Thus, $rpr(patlen)$ is simply the greatest k less than or equal to $patlen$ such that $k \leq 1$ or $pat(k - 1) \neq pat(patlen)$.

Of course we do not actually have two versions of $delta_1$. Instead we use only $delta_0$, and in place of $delta_1$ in the *max* expression we merely use the $delta_0$ entry unless it is *large* (in which case we use 0).

Note that the *fast* loop just scans down *string*, effectively looking for the last character $pat(patlen)$ in *pat*, skipping according to $delta_1$. ($delta_2$ can be ignored in this case since no terminal substring has yet been matched, i.e. $delta_2(patlen)$ is always less than or equal to the corresponding $delta_1$.) Control leaves this loop only when i exceeds *stringlen*. The test at *undo* decides whether this situation arose because all of *string* has been scanned or because $pat(patlen)$ was hit (which caused i to be incremented by *large*). If the first case obtains, *pat* does not occur in *string* and the algorithm returns false. If the second case obtains, then i is restored (by subtracting *large*) and we enter the *slow* loop which backs up checking for matches. When a mismatch is found we skip ahead by the maximum of the original $delta_1$ and $delta_2$ and reenter the *fast* loop. We estimate that 80 percent of the time spent in searching is spent in the *fast* loop.

The *fast* loop can be coded in four machine instructions:

```

fast: char ← string(i).
      i ← i + delta_0(char).
      skip the next instruction if i > stringlen.
      goto fast.
undo: . . .

```

We have implemented this algorithm in PDP-10 assembly language. In our implementation we have reduced the number of instructions in the *fast* loop to three by translating i down by *stringlen*; we can then test i against 0 and conditionally jump to *fast* in one instruction.

On a byte addressable machine it is easy to implement “ $char \leftarrow string(i)$ ” and “ $i \leftarrow i + delta_0(char)$ ” in one instruction each. Since our implementation was in PDP-10 assembly language we had to employ byte pointers to access characters in *string*. The PDP-10 instruction set provides an instruction for incrementing a byte pointer by one but not by other amounts. Our code therefore employs an array of 200 indexing byte pointers which we use to access characters in *string* in one indexed instruction (after computing the index) at the cost of a small (five-instruction) overhead every 200 characters. It should be noted that this trick only makes up for the lack of direct byte addressing; one can expect our algorithm to run somewhat faster on a byte-addressable machine.

6. Empirical Evidence

We have exhaustively tested the above PDP-10 implementation on random test data. To gather the test patterns we wrote a program which randomly selects a substring of a given length from a source string. We used this program to select 300 patterns of

length *patlen*, for each *patlen* from 1 to 14. We then used our algorithm to search for each of the test patterns in its source string, starting each search in a random position somewhere in the first half of the source string. All of the characters for both the patterns and the strings were in primary memory (rather than a secondary storage medium such as a disk).

We measured the cost of each search in two ways: the number of references made to *string* and the total number of machine instructions that actually got executed (ignoring the preprocessing to set up the two tables).

By dividing the number of references to *string* by the number of characters $i - 1$ passed before the pattern was found (or *string* was exhausted), we obtained the number of references to *string* per character passed. This measure is independent of the particular implementation of the algorithm. By dividing the number of instructions executed by $i - 1$, we obtained the average number of instructions spent on each character passed. This measure depends upon the implementation, but we feel that it is meaningful since the implementation is a straightforward encoding of the algorithm as described in the last section.

We then averaged these measures across all 300 samples for each pattern length.

Because the performance of the algorithm depends upon the statistical properties of *pat* and *string* (and hence upon the properties of the source string from which the test patterns were obtained), we performed this experiment for three different kinds of source strings, each of length 10,000. The first source string consisted of a random sequence of 0's and 1's. The second source string was a piece of English text obtained from an online manual. The third source string was a random sequence of characters from a 100-character alphabet.

In Figure 1 the average number of references to *string* per character in *string* passed is plotted against the pattern length for each of three source strings. Note that the number of references to *string* per character passed is less than 1. For example, for an English pattern of length 5, the algorithm typically inspects 0.24 characters for every character passed. That is, for every reference to *string* the algorithm passes about 4 characters, or, equivalently, the algorithm inspects only about a quarter of the characters it passes when searching for a pattern of length 5 in an English text string. Furthermore, the number of references per character drops as the patterns get longer. This evidence supports the conclusion that the algorithm is "sublinear" in the number of references to *string*.

For comparison, it should be noted that the Knuth, Morris, and Pratt algorithm references *string* precisely 1 time per character passed. The simple search algorithm references *string* about 1.1 times per character passed (determined empirically with the English sample above).

Fig. 1.





Fig. 2



In Figure 2 the average number of instructions executed per character passed is plotted against the pattern length. The most obvious feature to note is that the search speeds up as the patterns get longer. That is, the total number of instructions executed in order to pass over a character decreases as the length of the pattern increases.

Figure 2 also exhibits a second interesting feature of our implementation of the algorithm: For sufficiently

large alphabets and sufficiently long patterns the algorithm executes fewer than 1 instruction per character passed. For example, in the English sample, less than 1 instruction per character is executed for patterns of length 5 or more. Thus this implementation is “sub-linear” in the sense that it executes fewer than $i + patlen$ instructions before finding the pattern at i . This means that no algorithm which references each character it passes could possibly be faster than ours in these cases (assuming it takes at least one instruction to reference each character).

The best alternative algorithm for finding a single substring is that of Knuth, Morris, and Pratt. If that algorithm is implemented in the extraordinarily efficient way described in [4, pp. 11–12] and [2, Item 179],⁵ then the cost of looking at a character can be expected to be at least $3 - p$ instructions, where p is the probability that a character just fetched from *string* is equal to a given character of *pat*. Hence a horizontal line at $3 - p$ instructions/character represents the best (and, practically, the worst) the Knuth, Morris, and Pratt algorithm can achieve.

The simple string searching algorithm (when coded with a 3-instruction fast loop⁶) executes about 3.3 instructions per character (determined empirically on the English sample above).

As noted, the preprocessing time for our algorithm (and for Knuth, Morris, and Pratt) has been ignored. The cost of this preprocessing can be made linear in *patlen* (this is discussed further in the next section) and is trivial compared to a reasonably long search. We made no attempt to code this preprocessing efficiently. However, the average cost (in our implementation) ranges from 160 instructions (for strings of length 1) to about 500 instructions (for strings of length 14). It should be explained that our code uses a block transfer instruction to clear the 128-word δ_1 table at the beginning of the preprocessing, and we have counted this single instruction as though it were 128 instructions. This accounts for the unexpectedly large instruction count for preprocessing a one-character pattern.

7. Theoretical Analysis

The preprocessing for δ_1 requires an array the size of the alphabet. Our implementation first initializes all entries of this array to *patlen* and then sets up

⁵ This implementation automatically compiles *pat* into a machine code program which implicitly has the skip table built in and which is executed to perform the search itself. In [2] they compile code which uses the PDP-10 capability of fetching a character and incrementing a byte address in one instruction. This compiled code executes at least two or three instructions per character fetched from *string*, depending on the outcome of a comparison of the character to one from *pat*.

⁶ This loop avoids checking whether *string* is exhausted by assuming that the first character of *pat* occurs at the end of *string*. This can be arranged ahead of time. The loop actually uses the same three instruction codes used by the above-referenced implementation of the Knuth, Morris, and Pratt algorithm.

δ_1 in a linear scan through the pattern. Thus our preprocessing for δ_1 is linear in *patlen* plus the size of the alphabet.

At a slight loss of efficiency in the search speed one could eliminate the initialization of the δ_1 array by storing with each entry a key indicating the number of times the algorithm has previously been called. This approach still requires initializing the array the first time the algorithm is used.

To implement our algorithm for extremely large alphabets, one might implement the δ_1 table as a hash array. In the worst case, accessing δ_1 during the search itself could require order *patlen* instructions, significantly impairing the speed of the algorithm. Hence the algorithm as it stands almost certainly does not run in time linear in $i + patlen$ for infinite alphabets.

Knuth, in analyzing the algorithm, has shown that it still runs in linear time when δ_1 is omitted, and this result holds for infinite alphabets. Doing this, however, will drastically degrade the performance of the algorithm on the average. In [5] Knuth exhibits an algorithm for setting up δ_2 in time linear in *patlen*.

From the preceding empirical evidence, the reader can conclude that the algorithm is quite good in the average case. However, the question of its behavior in the worst case is nontrivial. Knuth has recently shed some light on this question. In [5] he proves that the execution of the algorithm (after preprocessing) is linear in $i + patlen$, assuming the availability of array space linear in *patlen* plus the size of the alphabet. In particular, he shows that in order to discover that *pat* does not occur in the first i characters of *string*, at most $6 * i$ characters from *string* are matched with characters in *pat*. He goes on to say that the constant 6 is probably much too large, and invites the reader to improve the theorem. His proof reveals that the linearity of the algorithm is entirely due to δ_2 .

We now analyze the average behavior of the algorithm by presenting a probabilistic model of its performance. As will become clear, the results of this analysis will support the empirical conclusions that the algorithm is usually “sublinear” both in the number of references to *string* and the number of instructions executed (for our implementation).

The analysis below is based on the following simplifying assumption: Each character of *pat* and *string* is an independent random variable. The probability that a character from *pat* or *string* is equal to a given character of the alphabet is p .

Imagine that we have just moved *pat* down *string* to a new position and that this position does not yield a match. We want to know the expected value of the ratio between the cost of discovering the mismatch and the distance we get to slide *pat* down upon finding the mismatch. If we define the cost to be the total number of references made to *string* before discovering the mismatch, we can obtain the expected value of the average number of references to *string* per character

passed. If we define the cost to be the total number of machine instructions executed in discovering the mismatch, we can obtain the expected value of the number of instructions executed per character passed.

In the following we say "only the last m characters of pat match" to mean "the last m characters of pat match the corresponding m characters in $string$ but the $(m + 1)$ -th character from the right end of pat fails to match the corresponding character in $string$."

The expected value of the ratio of cost to characters passed is given by:

$$\left(\sum_{m=0}^{patlen-1} cost(m) * prob(m) \right) / \left(\sum_{m=0}^{patlen-1} prob(m) * \left(\sum_{k=1}^{patlen} skip(m, k) * k \right) \right)$$

where $cost(m)$ is the cost associated with discovering that only the last m characters of pat match; $prob(m)$ is the probability that only the last m characters of pat match; and $skip(m, k)$ is the probability that, supposing only the last m characters of pat match, we will get to slide pat down by k .

Under our assumptions, the probability that only the last m characters of pat match is:

$$prob(m) = p^m(1 - p)/(1 - p^{patlen}).$$

(The denominator is due to the assumption that a mismatch exists.)

The probability that we will get to slide pat down by k is determined by analyzing how i is incremented. However, note that even though we increment i by the maximum max of the two $delta$ s, this will actually only slide pat down by $max - m$, since the increment of i also includes the m necessary to shift our attention back to the end of pat . Thus when we analyze the contributions of the two $delta$ s we speak of the amount by which they allow us to slide pat down, rather than the amount by which we increment i . Finally, recall that if the mismatched character $char$ occurs in the already matched final m characters of pat , then $delta_1$ is worthless and we always slide by $delta_2$. The probability that $delta_1$ is worthless is just $(1 - (1 - p)^m)$. Let us call this $probdelta_1worthless(m)$.

The conditions under which $delta_1$ will naturally let us slide forward by k can be broken down into four cases as follows: (a) $delta_1$ will let us slide down by 1 if $char$ is the $(m + 2)$ -th character from the righthand end of pat (or else there are no more characters in pat) and $char$ does not occur to the right of that position (which has probability $(1 - p)^m * (if\ m + 1 = patlen\ then\ 1\ else\ p)$). (b) $delta_1$ allows us to slide down k , where $1 < k < patlen - m$, provided the rightmost occurrence of $char$ in pat is $m + k$ characters from the right end of pat (which has probability $p * (1 - p)^{k+m-1}$). (c) When $patlen - m > 1$, $delta_1$ allows us to slide past $patlen - m$ characters if $char$ does not occur in pat at all (which has probability $(1 - p)^{patlen-1}$ given that we know $char$ is not the $(m + 1)$ -th character from

the right end of pat). Finally, (d) $delta_1$ never allows a slide longer than $patlen - m$ (since the maximum value of $delta_1$ is $patlen$).

Thus we can define the probability $probdelta_1(m, k)$ that when only the last m characters of pat match, $delta_1$ will allow us to move down by k as follows:

$$probdelta_1(m, k) = \begin{cases} \text{if } k = 1 \\ \text{then} \\ (1 - p)^m * (\text{if } m + 1 = patlen \text{ then } 1 \text{ else } p); \\ \text{elseif } 1 < k < patlen - m \text{ then } p * (1 - p)^{k+m-1}; \\ \text{elseif } k = patlen - m \text{ then } (1 - p)^{patlen-1}; \\ \text{else (i.e. } k > patlen - m) \text{ 0.} \end{cases}$$

(It should be noted that we will not put these formulas into closed form, but will simply evaluate them to verify the validity of our empirical evidence.)

We now perform a similar analysis for $delta_2$; $delta_2$ lets us slide down by k if (a) doing so sets up an alignment of the discovered occurrence of the last m characters of pat in $string$ with a plausible reoccurrence of those m characters elsewhere in pat , and (b) no smaller move will set up such an alignment. The probability $probpr(m, k)$ that the terminal substring of pat of length m has a plausible reoccurrence k characters to the left of its first character is:

$$probpr(m, k) = \begin{cases} \text{if } m + k < patlen \\ \text{then } (1 - p) * p^m \\ \text{else } p^{patlen-k} \end{cases}$$

Of course, k is just the distance $delta_2$ lets us slide provided there is no earlier reoccurrence. We can therefore define the probability $probdelta_2(m, k)$ that, when only the last m characters of pat match, $delta_2$ will allow us to move down by k recursively as follows:

$$probdelta_2(m, k) = probpr(m, k) \left(1 - \sum_{n=1}^{k-1} probdelta_2(m, n) \right).$$

We slide down by the maximum allowed by the two $delta$ s (taking adequate account of the possibility that $delta_1$ is worthless). If the values of the $delta$ s were independent, the probability that we would actually slide down by k would just be the sum of the products of the probabilities that one of the $delta$ s allows a move of k while the other allows a move of less than or equal to k .

However, the two moves are not entirely independent. In particular, consider the possibility that $delta_1$ is worthless. Then the $char$ just fetched occurs in the last m characters of pat and does not match the $(m + 1)$ -th. But if $delta_2$ gives a slide of 1 it means that sliding these m characters to the left by 1 produces a match. This implies that all of the last m characters of pat are equal to the character $m + 1$ from the right. But this character is known not to be $char$. Thus $char$ cannot occur in the last m characters of pat , violating the hypothesis that $delta_1$ was worthless. Therefore if $delta_1$ is worthless, the probability that $delta_2$ specifies a skip of 1 is 0 and the probability that it specifies one of the larger skips is correspondingly increased.

This interaction between the two δ 's is also felt (to a lesser extent) for the next m possible δ_2 's, but we ignore these (and in so doing accept that our analysis may predict slightly worse results than might be expected since we allow some short δ_2 moves when longer ones would actually occur).

The probability that δ_2 will allow us to slide down by k when only the last m characters of pat match, assuming that δ_1 is worthless, is:

$$\begin{aligned}
 \text{probdelta}'_2(m, k) &= \text{if } k = 1 \\
 &\quad \text{then } 0 \\
 &\quad \text{else} \\
 &\quad \text{probpr}(m, k) \left(1 - \sum_{n=2}^{k-1} \text{probdelta}'_2(m, n) \right).
 \end{aligned}$$

Finally, we can define $\text{skip}(m, k)$, the probability that we will slide down by k if only the last m characters of pat match:

$$\begin{aligned}
 \text{skip}(m, k) &= \text{if } k = 1 \\
 &\quad \text{then } \text{probdelta}_1(m, 1) * \text{probdelta}_2(m, 1) \\
 &\quad \text{else } \text{probdelta}_{1, \text{worthless}}(m) * \text{probdelta}'_2(m, k) \\
 &\quad + \sum_{n=1}^{k-1} \text{probdelta}_1(m, k) * \text{probdelta}_2(m, n) \\
 &\quad + \sum_{n=1}^{k-1} \text{probdelta}_1(m, n) * \text{probdelta}_2(m, k) \\
 &\quad + \text{probdelta}_1(m, k) * \text{probdelta}_2(m, k).
 \end{aligned}$$

Now let us consider the two alternative cost functions. In order to analyze the number of references to string per character passed over, $\text{cost}(m)$ should just be $m + 1$, the number of references necessary to confirm that only the last m characters of pat match.

In order to analyze the number of instructions executed per character passed over, $\text{cost}(m)$ should be the total number of instructions executed in discovering that only the last m characters of pat match. By inspection of our PDP-10 code:

$$\text{cost}(m) = \text{if } m = 0 \text{ then } 3 \text{ else } 12 + 6m.$$

We have computed the expected value of the ratio of cost per character skipped by using the above formulas (and both definitions of cost). We did so for pattern lengths running from 1 to 14 (as in our empirical evidence) and for the values of p appropriate for the three source strings used: For a random binary string p is 0.5, for an arbitrary English string it is (approximately) 0.09, and for a random string over a 100-character alphabet it is 0.01. The value of p for English was determined using a standard frequency count for the alphabetic characters [3] and empirically determining the frequency of space, carriage return, and line feed to be 0.23, 0.03, and 0.03, respectively.⁷

In Figure 3 we have plotted the theoretical ratio of references to string per character passed over against

⁷ We have determined empirically that the algorithm's performance on truly random strings where $p = 0.09$ is virtually identical to its performance on English strings. In particular, the reference count and instruction count curves generated by such random strings are almost coincidental with the English curves in Figures 1 and 2.

Fig. 3.





Fig. 4.



the pattern length. The most important fact to observe in Figure 3 is that the algorithm can be expected to make fewer than $i + \text{patlen}$ references to string before finding the pattern at location i . For example, for English text strings of length 5 or greater, the algorithm may be expected to make less than $(i + 5)/4$ references to string . The comparable figure for the Knuth,

Morris, and Pratt algorithm is of course precisely i . The figure for the intuitive search algorithm is always greater than or equal to i .

The reason the number of references per character passed decreases more slowly as $patlen$ increases is that for longer patterns the probability is higher that the character just fetched occurs somewhere in the pattern, and therefore the distance the pattern can be moved forward is shortened.

In Figure 4 we have plotted the theoretical ratio of the number of instructions executed per character passed versus the pattern length. Again we find that our implementation of the algorithm *can be expected* (for sufficiently large alphabets) to execute fewer than $i + patlen$ instructions before finding the pattern at location i . That is, our implementation is usually "sub-linear" even in the number of instructions executed. The comparable figure for the Knuth, Morris, and Pratt algorithm is at best $(3 - p) * (i + patlen - 1)$.⁸ For the simple search algorithm the expected value of the number of instructions executed per character passed is (approximately) 3.28 (for $p = 0.09$).

It is difficult to fully appreciate the role played by $delta_2$. For example, if the alphabet is large and patterns are short, then computing and trying to use $delta_2$ probably does not pay off much (because the chances are high that a given character in *string* does not occur anywhere in *pat* and one will almost always stay in the *fast* loop ignoring $delta_2$).⁹ Conversely, $delta_2$ becomes very important when the alphabet is small and the patterns are long (for now execution will frequently leave the *fast* loop; $delta_1$ will in general be small because many of the characters in the alphabet will occur in *pat* and only the terminal substring observations could cause large shifts). Despite the fact that it is difficult to appreciate the role of $delta_2$, it should be noted that the linearity result for the worst case behavior of the algorithm is due entirely to the presence of $delta_2$.

Comparing the empirical evidence (Figures 1 and 2) with the theoretical evidence (Figures 3 and 4, respectively), we note that the model is completely accurate for English and the 100-character alphabet. The model predicts much better behavior than we actually experience in the binary case. Our only explanation is that since $delta_2$ predominates in the binary alphabet and sets up alignments of the pattern and the string, the algorithm backs up over longer terminal substrings of the pattern before finding mismatches. Our analysis ignores this phenomenon.

⁸ Although the Knuth, Morris, and Pratt algorithm will fetch each of the first $i + patlen - 1$ characters of *string* precisely once, sometimes a character is involved in several tests against characters in *pat*. The number of such tests (each involving three instructions) is bounded by $\log_{\phi}(patlen)$, where ϕ is the golden ratio.

⁹ However, if the algorithm is implemented without $delta_2$, recall that, in exiting the *slow* loop, one must now take the *max* of $delta_1$ and $patlen - j + 1$ to allow for the possibility that $delta_1$ is worthless.

However, in summary, the theoretical analysis supports the conclusion that on the average the algorithm is sublinear in the number of references to *string* and, for sufficiently large alphabets and patterns, sublinear in the number of instructions executed (in our implementation).

8. Caveat Programmer

It should be observed that the preceding analysis has assumed that *string* is entirely in primary memory and that we can obtain the i th character in it in one instruction after computing its byte address. However, if *string* is actually on secondary storage, then the characters in it must be read in.¹⁰ This transfer will entail some time delay equivalent to the execution of, say, w instructions per character brought in, and (because of the nature of computer I/O) all of the first $i + patlen - 1$ characters will eventually be brought in whether we actually reference all of them or not. (A representative figure for w for paged transfers from a fast disk is 5 instructions/character.) Thus there may be a hidden cost of w instructions per character passed over.

According to the statistics presented above one might expect our algorithm to be approximately three times faster than the Knuth, Morris, and Pratt algorithm (for, say, English strings of length 6) since that algorithm executes about three instructions to our one. However, if the CPU is idle for the w instructions necessary to read each character, the actual ratios are closer to $w + 3$ instructions than to $w + 1$ instructions. Thus for paged disk transfers our algorithm can only be expected to be roughly $4/3$ faster (i.e. $5 + 3$ instructions to $5 + 1$ instructions) if we assume that we are idle during I/O. Thus for large values of w the difference between the various algorithms diminishes if the CPU is idle during I/O.

Of course, in general, programmers (or operating systems) try to avoid the situation in which the CPU is idle while awaiting an I/O transfer by overlapping I/O with some other computation. In this situation, the chances are that our algorithm will be I/O bound (we will search a page faster than it can be brought in), and indeed so will that of Knuth, Morris, and Pratt if $w > 3$. Our algorithm will require that fewer CPU cycles be devoted to the search itself so that if there are other jobs to perform, there will still be an overall advantage in using the algorithm.

¹⁰ We have implemented a version of our algorithm for searching through disk files. It is available as the subroutine FFILEPOS in the latest release of INTERLISP-10. This function uses the TENEX page mapping capability to identify one file page at a time with a buffer area in virtual memory. In addition to being faster than reading the page by conventional methods, this means the operating system's memory management takes care of references to pages which happen to still be in memory, etc. The algorithm is as much as 50 times faster than the standard INTERLISP-10 FILEPOS function (depending on the length of the pattern).

There are several situations in which it may not be advisable to use our algorithm. If the expected penetration i at which the pattern is found is small, the preprocessing time is significant and one might therefore consider using the obvious intuitive algorithm.

As previously noted, our algorithm can be most efficiently implemented on a byte-addressable machine. On a machine that does not allow byte addresses to be incremented and decremented directly, two possible sources of inefficiency must be addressed: The algorithm typically skips through *string* in steps larger than 1, and the algorithm may back up through *string*. Unless these processes are coded efficiently, it is probably not worthwhile to use our algorithm.

Furthermore, it should be noted that because the algorithm can back up through *string*, it is possible to cross a page boundary more than once. We have not found this to be a serious source of inefficiency. However, it does require a certain amount of code to handle the necessary buffering (if page I/O is being handled directly as in our FFILEPOS). One beauty of the Knuth, Morris, and Pratt algorithm is that it avoids this problem altogether.

A final situation in which it is unadvisable to use our algorithm is if the string matching problem to be solved is actually more complicated than merely finding the first occurrence of a single substring. For example, if the problem is to find the first of several possible substrings or to identify a location in *string* defined by a regular expression, it is much more advantageous to use an algorithm such as that of Aho and Corasick [1].

It may of course be possible to design an algorithm that searches for multiple patterns or instances of regular expressions by using the idea of starting the match at the right end of the pattern. However, we have not designed such an algorithm.

9. Historical Remarks

Our earliest formulation of the algorithm involved only δ_1 and implemented Observations 1, 2, and 3(a). We were aware that we could do something along the lines of δ_2 and Observation 3(b), but did not precisely formulate it. Instead, in April 1974, we coded the δ_1 version of the algorithm in Interlisp, merely to test its speed. We considered coding the algorithm in PDP-10 assembly language but abandoned the idea as impractical because of the cost of incrementing byte pointers by arbitrary amounts.

We have since learned that R.W. Gosper, of Stanford University, simultaneously and independently discovered the δ_1 version of the algorithm (private communication).

In April 1975, we started thinking about the implementation again and discovered a way to increment byte pointers by indexing through a table. We then formulated a version of δ_2 and coded the algorithm

more or less as it is presented here. This original definition of δ_2 differed from the current one in the following respect: If only the last m characters of *pat* (call this substring *subpat*) were matched, δ_2 specified a slide to the second from the rightmost occurrence of *subpat* in *pat* (allowing this occurrence to “fall off” the left end of *pat*) but without any special consideration of the character preceding this occurrence.

The average behavior of that version of the algorithm was virtually indistinguishable from that presented in this paper for large alphabets, but was somewhat worse for small alphabets. However, its worst case behavior was quadratic (i.e. required on the order of $i * patlen$ comparisons). For example, consider searching for a pattern of the form $CA(BA)^r$ in a string of the form $((XX)^r(AA)(BA)^r)^*$ (e.g. $r = 2$, *pat* = “CABABA,” and *string* = “XXXXAABA-BAXXXAABABA . . .”). The original definition of δ_2 allowed only a slide of 2 if the last “BA” of *pat* was matched before the next “A” failed to match. Of course in this situation this only sets up another mismatch at the same character in *string*, but the algorithm had to reinspect the previously inspected characters to discover it. The total number of references to *string* in passing i characters in this situation was $(r + 1) * (r + 2) * i / (4r + 2)$, where $r = (patlen - 2) / 2$. Thus the number of references was on the order of $i * patlen$.

However, on the average the algorithm was blindly fast. To our surprise, it was several times faster than the string searching algorithm in the Tenex TECO text editor. This algorithm is reputed to be quite an efficient implementation of the simple search algorithm because it searches for the first character of *pat* one full word at a time (rather than one byte at a time).

In the summer of 1975, we wrote a brief paper on the algorithm and distributed it on request.

In December 1975, Ben Kuipers of the M.I.T. Artificial Intelligence Laboratory read the paper and brought to our attention the improvement to δ_2 concerning the character preceding the terminal substring and its reoccurrence (private communication). Almost simultaneously, Donald Knuth of Stanford University suggested the same improvement and observed that the improved algorithm could certainly make no more than order $(i + patlen) * \log(patlen)$ references to *string* (private communication).

We mentioned this improvement in the next revision of the paper and suggested an additional improvement, namely the replacement of both δ_1 and δ_2 by a single two-dimensional table. Given the mismatched *char* from *string* and the position j in *pat* at which the mismatch occurred, this table indicated the distance to the last occurrence (if any) of the substring $[char, pat(j + 1), \dots, pat(patlen)]$ in *pat*. The revised paper concluded with the question of whether this improvement or a similar one produced an algorithm

which was at worst linear and on the average "sub-linear."

In January 1976, Knuth [5] proved that the simpler improvement in fact produces linear behavior, even in the worst case. We therefore revised the paper again and gave δ_2 its current definition.

In April 1976, R.W. Floyd of Stanford University discovered a serious statistical fallacy in the first version of our formula giving the expected value of the ratio of cost to characters passed. He provided us (private communication) with the current version of this formula.

Thomas Standish, of the University of California at Irvine, has suggested (private communication) that the implementation of the algorithm can be improved by fetching larger bytes in the *fast* loop (i.e. bytes containing several characters) and using a hash array to encode the extended δ_1 table. Provided the difficulties at the boundaries of the pattern are handled efficiently, this could improve the behavior of the algorithm enormously since it exponentially increases the effective size of the alphabet and reduces the frequency of common characters.

Acknowledgments. We would like to thank B. Kuipers, of the M.I.T. Artificial Intelligence Labora-

tory, for his suggestion concerning δ_2 and D. Knuth, of Stanford University, for his analysis of the improved algorithm. We are grateful to the anonymous reviewer for *Communications* who suggested the inclusion of evidence comparing our algorithm with that of Knuth, Morris, and Pratt, and for the warnings contained in Section 8. B. Mont-Reynaud, of the Stanford Research Institute, and L. Guibas, of Xerox Palo Alto Research Center, proofread drafts of this paper and suggested several clarifications. We would also like to thank E. Taft and E. Fiala of Xerox Palo Alto Research Center for their advice regarding machine coding the algorithm.

Received June 1975; revised April 1976

References

1. Aho, A.V., and Corasick, M.J. Fast pattern matching: An aid to bibliographic search. *Comm. ACM* 18, 6 (June, 1975), 333-340.
2. Beeler, M., Gosper, R.W., and Schroepel, R. Hakmem. Memo No. 239, M.I.T. Artificial Intelligence Lab., M.I.T., Cambridge, Mass., Feb. 29, 1972.
3. Dewey, G. *Relativ Frequency of English Speech Sounds*. Harvard U. Press, Cambridge, Mass., 1923, p. 185.
4. Knuth, D.E., Morris, J.H., and Pratt, V.R. Fast pattern matching in strings. TR CS-74-440, Stanford U., Stanford, Calif., 1974.
5. Knuth, D.E., Morris, J.H., and Pratt, V.R. Fast pattern matching in strings. (to appear in *SIAM J. Comput.*)

Professional Activities Calendar of Events

ACM's calendar policy is to list open computer science meetings that are held on a not-for-profit basis. Not included in the calendar are educational seminars, institutes, and courses. Submittals should be substantiated with name of the sponsoring organization, fee schedule, and chairman's name and full address.

One telephone number contact for those interested in attending a meeting will be given when a number is specified for this purpose in the news release text or in a direct communication to this periodical.

All requests for ACM sponsorship or cooperation should be addressed to Chairman, Conferences and Symposia Committee, Dr. W.S. Dorsey, Dept. 503/504 Rockwell International Corporation, Anaheim, CA 92803. For European events, a copy of the request should also be sent to the European Regional Representative. Technical Meeting Request Forms for this purpose can be obtained from ACM Headquarters or from the European Regional Representative. Lead time should include 2 months (3 months if for Europe) for processing of the request, plus the necessary months (minimum 2) for any publicity to appear in *Communications*.

Events for which ACM or a subunit of ACM is a sponsor or collaborator are indicated by ■. Dates precede titles.

In this issue the calendar is given to April 1978. New Listings are shown first; they will appear next month as Previous Listings.

NEW LISTINGS

16-17 November 1977
■ **Workshop on Future Directions in Computer Architecture**, Austin, Tex. Sponsors: ACM SIG-ARCH, IEEE-CS TCCA, and University of Texas at Austin. Conf. chm: G. Jack Lipovski, Dept. of EE, University of Texas, Austin, TX 78712.
5-9 December 1977

■ **Third International Symposium on Computing Methods in Applied Sciences and Engineering**, Versailles, France. Organized by IRIA. Sponsors: AFCET, GAMNI, IFIP WG7.2 Contact: Institut de Recherche D'Informatique et D'Automatique, Domaine de Voluceau, Rocquencourt, 78150 Le Chesnay, France.

13-15 February 1978
■ **Symposium on Computer Network Protocols**, Liège, Belgium. Sponsors: IFIP T.C.6 and ACM Belgian Chapter. Contact: A. Danthine, Symposium on Computer Network Protocols, Avenue des Tilleuls, 49, B-4000, Liège, Belgium.
3 March 1978

■ **Indiana University Computer Network Conference on Instructional Computing Applications**, Indiana University-East, Richmond, Ind. Sponsor:

Indiana University Computing Network. Chm: Tom Osgood, IU-EAST, 2325 Chester Boulevard, Richmond, IN 47374.

12-17 March 1978
■ **Symposium on Computer Simulation of Bulk Matter from Molecular Perspective**, Anaheim, Calif.; part of 1978 Annual Spring Meeting of American Chemical Society. Sponsor: ACS Div. Computers in Chemistry. Contact: Peter Lykos, Illinois Institute of Technology, Chicago, IL 60616; 312 567-3430.

28-30 March 1978
■ **3rd Symposium on Programming**, Paris, France. Sponsor: Centre National de la Recherche Scientifique (CNRS) and Université Pierre et Marie Curie. Contact: Secrétariat du Colloque, Institut de Programmation, 4, Place Jussieu, 75230 Paris Cedex 05, France.

29-31 March 1978
■ **Conference on Information Sciences and Systems**, Johns Hopkins University, Baltimore, Md. Contact: 1978 CISS, Department of Electrical Engineering, Johns Hopkins University, Baltimore, MD 21218.

3-7 April 1978
■ **Fifth International Symposium on Computing in Literary and Linguistic Research**, University of Aston, Birmingham, England. Sponsor: Association for Literary and Linguistic Computing. Contact: The Secretary (CLLR), Modern Languages Dept., University of Aston, Birmingham B4 7ET, England.

4-8 April 1978
■ **Second International Conference on Combinatorial Mathematics**, Barbizon-Plaza Hotel, New York City. Sponsor: New York Academy of Sciences. Contact: Conference Dept., New York Academy of Sciences, 2 East 63 St., New York, NY 10021; 212 838-0230.

15-19 May 1978
■ **16th Annual Convention of the Association for Educational Data Systems**, Atlanta, Ga. Sponsor: AEDS. Contact: James E. Eisele, Office of Computing Activities, University of Georgia, Athens, GA 30602.

22-25 May 1978
■ **Sixth International CODATA Conference**, Taormina, Italy. Sponsor: International Council of Scientific Unions Comm. on Data for Science and Technology. Contact: CODATA Secretariat, 51, Boulevard de Montmorency, 75016 Paris, France.

24-26 May 1978
■ **1978 SIAM National Meeting**, University of Wisconsin, Madison, Wis. Sponsor: SIAM in cooperation with ACM SIGSAM. Contact: H.B. Hair, SIAM, 33 South 17 St., Philadelphia, PA 19103; 215 564-2929.

26 May 1978
■ **Computer Algebra Symposium**, University of Wisconsin, Madison, Wis.; part of the 1978 SIAM National Meeting. Sponsor: SIAM in cooperation with ACM SIGSAM. Symp. chm: George E.

Collins, Computer Sciences Dept., University of Wisconsin, 1210 W. Dayton Street, Madison WI 53706.

12-16 June 1978
■ **7th Triennial IFAC World Congress**. Sponsor: IFAC. Contact: IFAC 78 Secretariat, POB 192, 00101 Helsinki 10, Finland.

19-22 June 1978
■ **Annual Conference of the American Society for Engineering Education** (Computers in Education Division Program), University of British Columbia, Vancouver, B.C., Canada. Sponsor: ASEE Computers in Education Division. Contact: ASEE, Suite 400, One DuPont Circle, Washington, DC 20036.

22-23 June 1978
■ **International Conference on the Performance of Computer Installations**, Gardone Riviera, Lake Garda, Italy. Sponsor: Sperry Univac, Italy, with cooperation of ACM SIGMETRICS, ECOMA, AICA, ACM Italian Chapter. Contact: Conference Secretariat, CILEA, Via Raffaello Sanzio 4, 20090 Segrate, Milan, Italy.

2-4 August 1978
■ **International Conference on Databases: Improving Usability and Responsiveness**, Technion, Haifa, Israel. Sponsor: Technion in cooperation with ACM. Prog. chm: Ben Shneiderman, Dept. of Information Systems Management, University of Maryland, College Park, MD 20742.

13-18 August 1978
■ **Symposium on Modeling and Simulation Methodology**, Weizmann Institute of Science, Rehovot, Israel. Contact: H.J. Highland, State University Technical College Farmingdale, N.Y., or B.P. Zeigler, Dept. of Applied Mathematics, Weizmann Institute of Science, Rehovot, Israel.
30 October-1 November 1978

■ **1978 SIAM Fall Meeting**, Hyatt Regency Hotel, Knoxville, Tenn. Sponsor: SIAM. Contact: H.B. Hair, SIAM, 33 South 17th St., Philadelphia, PA 19103; 215 564-2929.

PREVIOUS LISTINGS

17-19 October 1977
■ **ACM 77 Annual Conference**, Olympic Hotel, Seattle, Wash. Gen. chm: James S. Ketchel, Box 16156, Seattle, WA 98116; 206 935-6776.

17-21 October 1977
■ **Systems 77, Computer Systems and Their Application**, Munich, Federal Republic of Germany. Contact: Münchener Messe- und Ausstellungensgesellschaft mbh, Kongresszentrum, Kongressbüro Systems 77, Postfach 12 10 09, D-8000 München 12, Federal Republic of Germany.

18-19 October 1977
■ **MSFC/VAH Data Management Symposium**, Sheraton Motor Inn, Huntsville, Ala. Sponsors: NASA Marshall Space Flight Center, University of Alabama in Huntsville. Contact: General Chair
(Calendar continued on p. 781)