

linearization of the circular string. If $l = 0$ or $l = n + 1$, then cut the circular string between character n and character 1. Each leaf in the subtree of this point gives a cutting point yielding the same linear string.

The correctness of this solution is easy to establish and is left as an exercise.

This method runs in linear time and is therefore time optimal. A different linear-time method with a smaller constant was given by Shiloach [404].

7.14. APL13: Suffix arrays – more space reduction

In Section 6.5.1, we saw that when alphabet size is included in the time and space bounds, the suffix tree for a string of length m either requires $\Theta(m|\Sigma|)$ space or the minimum of $O(m \log m)$ and $O(m \log |\Sigma|)$ time. Similarly, searching for a pattern P of length n using a suffix tree can be done in $O(n)$ time only if $\Theta(m|\Sigma|)$ space is used for the tree, or if we assume that up to $|\Sigma|$ character comparisons cost only constant time. Otherwise, the search takes the minimum of $O(n \log m)$ and $O(n \log |\Sigma|)$ comparisons. For these reasons, a suffix tree may require too much space to be practical in some applications. Hence a more space efficient approach is desired that still retains most of the advantages of searching with a suffix tree.

In the context of the substring problem (see Section 7.3) where a fixed string T will be searched many times, the key issues are the time needed for the search and the space used by the fixed data structure representing T . The space used during the preprocessing of T is of less concern, although it should still be “reasonable”.

Manber and Myers [308] proposed a new data structure, called a *suffix array*, that is very space efficient and yet can be used to solve the exact matching problem or the substring problem almost as efficiently as with a suffix tree. Suffix arrays are likely to be an important contribution to certain string problems in computational molecular biology, where the alphabet can be large (we will discuss some of the reasons for large alphabets below). Interestingly, although the more formal notion of a suffix array and the basic algorithms for building and using it were developed in [308], many of the ideas were anticipated in the biological literature by Martinez [310].

After defining suffix arrays we show how to convert a suffix tree to a suffix array in linear time. It is important to be clear on the setting of the problem. String T will be held fixed for a long time, while P will vary. Therefore, the goal is to find a space-efficient representation for T (a suffix array) that will be held fixed and that facilitates search problems in T . However, the amount of space used during the construction of that representation is not so critical. In the exercises we consider a more space efficient way to build the representation itself.

Definition Given an m -character string T , a *suffix array* for T , called Pos , is an array of the integers in the range 1 to m , specifying the lexicographic order of the m suffixes of string T .

That is, the suffix starting at position $Pos(1)$ of T is the lexically smallest suffix, and in general suffix $Pos(i)$ of T is lexically smaller than suffix $Pos(i + 1)$.

As usual, we will affix a terminal symbol $\$$ to the end of S , but now we interpret it to be lexically *less* than any other character in the alphabet. This is in contrast to its interpretation in the previous section. As an example of a suffix array, if T is *mississippi*, then the suffix array Pos is 11, 8, 5, 2, 1, 10, 9, 7, 4, 6, 3. Figure 7.4 lists the eleven suffixes in lexicographic order.

```

11: i
8: ippi
5: issippi
2: ississippi
1: mississippi
10: pi
9: ppi
7: sippi
4: sisippi
6: ssippi
3: ssissippi

```

Figure 7.4: The eleven suffixes of *mississippi* listed in lexicographic order. The starting positions of those suffixes define the suffix array *Pos*.

Notice that the suffix array holds only integers and hence contains no information about the alphabet used in string T . Therefore, the space required by suffix arrays is modest – for a string of length m , the array can be stored in exactly m computer words, assuming a word size of at least $\log m$ bits.

When augmented with an additional $2m$ values (called *Lcp* values and defined later), the suffix array can be used to find all the occurrences in T of a pattern P in $O(n + \log_2 m)$ single-character comparison and bookkeeping operations. Moreover, this bound is independent of the alphabet size. Since for most problems of interest $\log_2 m$ is $O(n)$, the substring problem is solved by using suffix arrays as efficiently as by using suffix trees.

7.14.1. Suffix tree to suffix array in linear time

We assume that sufficient space is available to build a suffix tree for T (this is done once during a preprocessing phase), but that the suffix tree cannot be kept intact to be used in the (many) subsequent searches for patterns in T . Instead, we convert the suffix tree to the more space efficient suffix array. Exercises 53, 54, and 55 develop an alternative, more space efficient (but slower) method, for *building* a suffix array.

A suffix array for T can be obtained from the suffix tree \mathcal{T} for T by performing a “lexical” depth-first traversal of \mathcal{T} . Once the suffix array is built, the suffix tree is discarded.

Definition Define an edge (v, u) to be *lexically less* than an edge (v, w) if and only if the first character on the (v, u) edge is lexicographically less than the first character on (v, w) . (In this application, the end of string character $\$$ is lexicographically less than any other character.)

Since no two edges out of v have labels beginning with the same character, there is a strict lexical ordering of the edges out of v . This ordering implies that the path from the root of \mathcal{T} following the lexicographically smallest edge out of each encountered node leads to a leaf of \mathcal{T} representing the lexicographically smallest suffix of T . More generally, a depth-first traversal of \mathcal{T} that traverses the edges out of each node v in their lexical order will encounter the leaves of \mathcal{T} in the lexical order of the suffixes they represent. Suffix array *Pos* is therefore just the ordered list of suffix numbers encountered at the leaves of \mathcal{T} during the lexical depth-first search. The suffix tree for T is constructed in linear time, and the traversal also takes only linear time, so we have the following:

Theorem 7.14.1. *The suffix array *Pos* for a string T of length m can be constructed in $O(m)$ time.*




Figure 7.5: The lexical depth-first traversal of the suffix tree visits the leaves in order 5, 2, 6, 3, 4, 1.

For example, the suffix tree for $T = \text{tartar}$ is shown in Figure 7.5. The lexical depth-first traversal visits the nodes in the order 5, 2, 6, 3, 4, 1, defining the values of array Pos .

As an implementation detail, if the branches out of each node of the tree are organized in a *sorted* linked list (as discussed in Section 6.5, page 116) then the overhead to do a lexical depth-first search is the same as for any depth-first search. Every time the search must choose an edge out of a node v to traverse, it simply picks the next edge on v 's linked list.

7.14.2. How to search for a pattern using a suffix array

The suffix array for string T allows a very simple algorithm to find all occurrences of any pattern P in T . The key is that if P occurs in T then all the locations of those occurrences will be grouped consecutively in Pos . For example, $P = \text{issi}$ occurs in mississippi starting at locations 2 and 5, which are indeed adjacent in Pos (see Figure 7.4). So to search for occurrences of P in T simply do binary search over the suffix array. In more detail, suppose that P is lexically less than the suffix in the middle position of Pos (i.e., suffix $Pos(\lceil m/2 \rceil)$). In that case, the first place in Pos that contains a position where P occurs in T must be in the first half of Pos . Similarly, if P is lexically greater than suffix $Pos(\lceil m/2 \rceil)$, then the places where P occurs in T must be in the second half of Pos . Using binary search, one can therefore find the smallest index i in Pos (if any) such that P exactly matches the first n characters of suffix $Pos(i)$. Similarly, one can find the largest index i' with that property. Then pattern P occurs in T starting at every location given by $Pos(i)$ through $Pos(i')$.

The lexical comparison of P to any suffix takes time proportional to the length of the common prefix of those two strings. That prefix has length at most n ; hence

Theorem 7.14.2. *By using binary search on array Pos , all the occurrences of P in T can be found in $O(n \log m)$ time.*

Of course, the true behavior of the algorithm depends on how many long prefixes of P occur in T . If very few long prefixes of P occur in T then it will rarely happen that a specific lexical comparison actually takes $\Theta(n)$ time and generally the $O(n \log m)$ bound is quite pessimistic. In “random” strings (even on large alphabets) this method should run in $O(n + \log m)$ expected time. In cases where many long prefixes of P do occur in T , then the method can be improved with the two tricks described in the next two subsections.

7.14.3. A simple accelerant

As the binary search proceeds, let L and R denote the left and right boundaries of the "current search interval". At the start, L equals 1 and R equals m . Then in each iteration of the binary search, a query is made at location $M = \lceil (R + L)/2 \rceil$ of Pos . The search algorithm keeps track of the longest prefixes of $Pos(L)$ and $Pos(R)$ that match a prefix of P . Let l and r denote those two prefix lengths, respectively, and let $mlr = \min(l, r)$.

The value mlr can be used to accelerate the lexical comparison of P and suffix $Pos(M)$. Since array Pos gives the lexical ordering of the suffixes of T , if i is any index between L and R , the first mlr characters of suffix $Pos(i)$ must be the same as the first mlr characters of suffix $Pos(L)$ and hence of P . Therefore, the lexical comparison of P and suffix $Pos(M)$ can begin from position $mlr + 1$ of the two strings, rather than starting from the first position.

Maintaining mlr during the binary search adds little additional overhead to the algorithm but avoids many redundant comparisons. At the start of the search, when $L = 1$ and $R = m$, explicitly compare P to suffix $Pos(1)$ and suffix $Pos(m)$ to find l, r , and mlr . However, the worst-case time for this revised method is still $O(n \log m)$. Myers and Manber report that the use of mlr alone allows the search to run as fast in practice as the $O(n + \log m)$ worst-case method that we first advertised. Still, if only because of its elegance, we present the full method that guarantees that better worst-case bound.

7.14.4. A super-accelerant

Call an examination of a character in P *redundant* if that character has been examined before. The goal of the acceleration is to reduce the number of redundant character examinations to at most one per iteration of the binary search – hence $O(\log m)$ in all. The desired time bound, $O(n + \log m)$, follows immediately. The use of mlr alone does not achieve this goal. Since mlr is the minimum of l and r , whenever $l \neq r$, all characters in P from $mlr + 1$ to the maximum of l and r will have already been examined. Thus any comparisons of those characters will be redundant. What is needed is a way to begin comparisons at the *maximum* of l and r .

Definition $Lcp(i, j)$ is the length of the longest common prefix of the suffixes specified in positions i and j of Pos . That is, $Lcp(i, j)$ is the length of the longest prefix common to suffix $Pos(i)$ and suffix $Pos(j)$. The term Lcp stands for *longest common prefix*.

For example, when $T = mississippi$, suffix $Pos(3)$ is *issippi*, suffix $Pos(4)$ is *ississippi*, and so $Lcp(3, 4)$ is four (see Figure 7.4).

To speed up the search, the algorithm uses $Lcp(L, M)$ and $Lcp(M, R)$ for each triple (L, M, R) that arises during the execution of the binary search. For now, we assume that these values can be obtained in constant time when needed and show how they help the search. Later we will show how to compute the particular Lcp values needed by the binary search during the preprocessing of T .

How to use Lcp values

Simplest case In any iteration of the binary search, if $l = r$, then compare P to suffix $Pos(M)$ starting from position $mlr + 1 = l + 1 = r + 1$, as before.

General case When $l \neq r$, let us assume without loss of generality that $l > r$. Then there are three subcases:




Figure 7.6: Subcase 1 of the super-accelerant. Pattern P is $abcdemn$, shown vertically running upwards from the first character. The suffixes $Pos(L)$, $Pos(M)$, and $Pos(R)$ are also shown vertically. In this case, $Lcp(L, M) > 0$ and $l > r$. Any starting location of P in T must occur in Pos to the right of M , since P agrees with suffix $Pos(M)$ only up to character l .

- If $Lcp(L, M) > l$, then the common prefix of suffix $Pos(L)$ and suffix $Pos(M)$ is longer than the common prefix of P and $Pos(L)$. Therefore, P agrees with suffix $Pos(M)$ up through character l . In other words, characters $l + 1$ of suffix $Pos(L)$ and suffix $Pos(M)$ are identical and lexically less than character $l + 1$ of P (the last fact follows since P is lexically greater than suffix $Pos(L)$). Hence all (if any) starting locations of P in T must occur to the right of position M in Pos . So in any iteration of the binary search where this case occurs, no examinations of P are needed; L just gets changed to M , and l and r remain unchanged. (See Figure 7.6.)
- If $Lcp(L, M) < l$, then the common prefix of suffix $Pos(L)$ and $Pos(M)$ is smaller than the common prefix of suffix $Pos(L)$ and P . Therefore, P agrees with suffix $Pos(M)$ up through character $Lcp(L, M)$. The $Lcp(L, M) + 1$ characters of P and suffix $Pos(L)$ are identical and lexically less than character $Lcp(L, M) + 1$ of suffix $Pos(M)$. Hence all (if any) starting locations of P in T must occur to the left of position M in Pos . So in any iteration of the binary search where this case occurs, no examinations of P are needed; r is changed to $Lcp(L, M)$, l remains unchanged, and R is changed to M .
- If $Lcp(L, M) = l$, then P agrees with suffix $Pos(M)$ up to character l . The algorithm then lexically compares P to suffix $Pos(M)$ starting from position $l + 1$. In the usual manner, the outcome of that lexical comparison determines which of L or R change, along with the corresponding change of l or r .

Theorem 7.14.3. *Using the Lcp values, the search algorithm does at most $O(n + \log m)$ comparisons and runs in that time.*

PROOF First, by simple case analysis it is easy to verify that neither l nor r ever decrease during the binary search. Also, every iteration of the binary search terminates the search, examines no characters of P , or ends after the first mismatch occurs in that iteration.

In the two cases ($l = r$ or $Lcp(L, M) = l > r$) where the algorithm examines a character during the iteration, the comparisons start with character $\max(l, r)$ of P . Suppose there are k characters of P examined in that iteration. Then there are $k - 1$ matches during the iteration, and at the end of the iteration $\max(l, r)$ increases by $k - 1$ (either l or r is changed to that value). Hence at the start of any iteration, character $\max(l, r)$ of P may have already been examined, but the next character in P has not been. That means at most one redundant comparison per iteration is done. Thus no more than $\log_2 m$ redundant comparisons are done overall. There are at most n nonredundant comparisons of characters




Figure 7.7: Binary tree B representing all the possible search intervals in any execution of binary search in a list of length $m = 8$.

of P , giving a total bound of $n + \log m$ comparisons. All the other work in the algorithm can clearly be done in time proportional to these comparisons. \square

7.14.5. How to obtain the Lcp values

The Lcp values needed to accelerate searches are precomputed in the preprocessing phase during the creation of the suffix array. We first consider how many possible Lcp values are ever needed (over any possible execution of binary search). For convenience, assume m is a power of two.

Definition Let B be a complete binary tree with m leaves, where each node of B is labeled with a pair of integers (i, j) , $1 \leq i \leq j \leq m$. The root of B is labeled $(1, m)$. Every nonleaf node (i, j) has two children; the left one is labeled $(i, \lfloor (i+j)/2 \rfloor)$, and the right one is labeled $(\lfloor (i+j)/2 \rfloor, j)$. The leaves of B are labeled $(i, i+1)$ (plus one labeled $(1, 1)$) and are ordered left to right in increasing order of i . (See Figure 7.7.)

Essentially, the node labels specify the endpoints (L, R) of all the possible search intervals that could arise in the binary search of an ordered list of length m . Since B is a binary tree with m leaves, B has $2m - 1$ nodes in total. So there are only $O(m)$ Lcp values that need be precomputed. It is therefore plausible that those values can be accumulated during the $O(m)$ -time preprocessing of T ; but how exactly? In the next lemma we show that the Lcp values at the leaves of B are easy to accumulate during the lexical depth-first traversal of T .

Lemma 7.14.1. *In the depth-first traversal of T , consider the internal nodes visited between the visits to leaf $Pos(i)$ and leaf $Pos(i+1)$, that is, between the i th leaf visited and the next leaf visited. From among those internal nodes, let v denote the one that is closest to the root. Then $Lcp(i, i+1)$ equals the string-depth of node v .*

For example, consider again the suffix tree shown in Figure 7.5 (page 151). $Lcp(5,6)$ is the string-depth of the parent of leaves 4 and 1. That string-depth is 3, since the parent of 4 and 1 is labeled with the string tar . The values of $Lcp(i, i+1)$ are 2, 0, 1, 0, 3 for i from 1 to 5.

The hardest part of Lemma 7.14.1 involves parsing it. Once done, the proof is immediate from properties of suffix trees, and it is left to the reader.