

AC Automaton for the Set of Suffixes

As already mentioned, a suffix tree with suffix links is essentially an Aho–Corasick automaton for the set of all suffixes.

- We saw that it is possible to follow suffix link / failure transition from any locus, not just from suffix tree nodes.
- Following such an implicit suffix link may take more than a constant time, but the total time during the scanning of a string with the automaton is linear in the length of the string. This can be shown with a similar argument as in the construction algorithm.

Thus suffix tree is asymptotically as fast to operate as the AC automaton, but needs much less space.

177

Longest Common Extension

The **longest common extension** (LCE) query asks for the length of the **longest common prefix** of two suffixes of a text T :

$$LCE(i, j) := lcp(T_i, T_j)$$

- The **lowest common ancestor** (LCA) of two nodes u and v in a tree is the deepest node that is an ancestor of both u and v . Any tree can be preprocessed in **linear time** so that the LCA of any two nodes can be computed in **constant time**. The details are omitted here.
- A LCE query can be implemented as a LCA query on the suffix tree of T :

$$LCE(i, j) = LCA(w_i, w_j)$$

where w_i and w_j are the leaves that represent the suffixes T_i and T_j . Thus, given the suffix tree augmented with a precomputed LCA data structure, LCE queries can be answered in constant time.

Some $\mathcal{O}(kn)$ worst case time approximate string matching algorithms are based on LCE queries.

179

Suffix Array

The suffix array of a text T is a lexicographically ordered array of the set $T_{[0..n]}$ of all suffixes of T . More precisely, the suffix array is an array $SA[0..n]$ of integers containing a permutation of the set $[0..n]$ such that $T_{SA[0]} < T_{SA[1]} < \dots < T_{SA[n]}$.

A related array is the **inverse suffix array** SA^{-1} which is the inverse permutation, i.e., $SA^{-1}[SA[i]] = i$ for all $i \in [0..n]$. The value $SA^{-1}[j]$ is the **lexicographical rank** of the suffix T_j .

As with suffix trees, it is common to add the end symbol $T[n] = \$$. It has no effect on the suffix array assuming $\$$ is smaller than any other symbol.

Example 4.7: The suffix array and the inverse suffix array of the text $T = \text{banana}\$$.

i	$SA[i]$	$T_{SA[i]}$	j	$SA^{-1}[j]$	T_j
0	6	\$	0	4	banana\$
1	5	a\$	1	3	anana\$
2	3	ana\$	2	6	nana\$
3	1	anana\$	3	2	ana\$
4	0	banana\$	4	5	na\$
5	4	na\$	5	1	a\$
6	2	nana\$	6	0	\$

181

Exact String Matching

As with suffix trees, **exact string matching** in T can be performed by a **prefix search** on the suffix array. The answer can be conveniently given as a **contiguous interval** $SA[b..e]$ that contains the suffixes with the given prefix. The interval can be found using string binary search.

- If we have the additional arrays $LLCP$ and $RLCP$, the result interval can be computed in $\mathcal{O}(|P| + \log n)$ time.
- Without the additional arrays, we have $\mathcal{O}(|P| + \log n)$ average time complexity, and we can achieve $\mathcal{O}(|P| \log_{|P|} n)$ worst case time with the skewed string binary search (Algorithm 1.40), and even better with a more complicated algorithm (see slide 59).
- We can then count the number of occurrences in $\mathcal{O}(1)$ time, list all occ occurrences in $\mathcal{O}(occ)$ time, or list a sample of k occurrences in $\mathcal{O}(k)$ time.

An alternative algorithm for computing the interval $SA[b..e]$ is called **backward search**. It is commonly used with compressed representations of suffix arrays.

183

Matching Statistics

The matching statistics of a string $S[0..n]$ with respect to a string T is an array $MS[0..n]$, where $MS[i]$ is a pair (ℓ_i, p_i) such that

1. $S[i..i + \ell_i]$ is the longest prefix of S_i that is a factor of T , and
2. $T[p_i..p_i + \ell_i] = S[i..i + \ell_i]$.

Matching statistics can be computed by using the suffix tree of T as an AC-automaton and scanning S with it.

- If before reading $S[i]$ we are at the locus (v, d) in the automaton, then $S[i - d..i] = T[j..j + d]$, where $j = \text{start}(v)$. If reading $S[i]$ causes a failure transition, then $MS[i - d] = (d, j)$.
- Following the failure transition decrements d and thus increments $i - d$ by one. Following a normal transition/edge, increments both i and d by one, and thus $i - d$ stays the same. Thus all entries are computed.

From the matching statistics, we can easily compute the longest common factor of S and T . Because we need the suffix tree only for T , this saves space compared to a generalized suffix tree.

Matching statistics are also used in some approximate string matching algorithms.

178

Longest Palindrome

A palindrome is a string that is its own reverse. For example, **saippuakauppias** is a palindrome.

We can use the LCA preprocessed generalized suffix tree of a string T and its reverse T^R to find the **longest palindrome** in T in linear time.

- Let k_i be the length of the longest common extension of T_{i+1} and T_{n-i}^R , which can be computed in constant time. Then $T[i - k_i..i + k_i]$ is the longest odd length palindrome with the middle at i .
- We can find the longest odd length palindrome by computing k_i for all $i \in [0..n]$ in $\mathcal{O}(n)$ time.
- The longest even length palindrome can be found similarly in $\mathcal{O}(n)$ time. The longest palindrome overall is the longer of the two.

180

Suffix array is much simpler data structure than suffix tree. In particular, the type and the size of the alphabet are usually not a concern.

- The size of the suffix array is $\mathcal{O}(n)$ on any alphabet.
- We will later see that the suffix array can be constructed in the same asymptotic time it takes to **sort the characters** of the text.

Suffix array construction algorithms are quite fast in practice too. Probably the fastest way to construct a suffix tree is to construct a suffix array first and then use it to construct the suffix tree. (We will see how in a moment.)

Suffix arrays are rarely used alone but are augmented with other arrays and data structures depending on the application. We will see some of them in the next slides.

182

LCP Array

Efficient string binary search uses the arrays $LLCP$ and $RLCP$. However, for many applications, the suffix array is augmented with the lcp array of Definition 1.11 (Lecture 2). For all $i \in [1..n]$, we store

$$LCP[i] = lcp(T_{SA[i]}, T_{SA[i-1]})$$

Example 4.8: The LCP array for $T = \text{banana}\$$.

i	$SA[i]$	$LCP[i]$	$T_{SA[i]}$
0	6		\$
1	5	0	a\$
2	3	1	ana\$
3	1	3	anana\$
4	0	0	banana\$
5	4	0	na\$
6	2	2	nana\$

184

Using the solution of Exercise 2.5 (construction of compact trie from sorted array and LCP array), the suffix tree can be constructed from the suffix and LCP arrays in linear time.

However, many suffix tree applications can be solved using the suffix and LCP arrays directly. For example:

- The **longest repeating factor** is marked by the maximum value in the LCP array.
- The **number of distinct factors** can be computed by the formula

$$\frac{n(n+1)}{2} + 1 - \sum_{i=1}^n LCP[i]$$


since it equals the number of nodes in the uncompact suffix trie, for which we can use Theorem 1.17.

- **Matching statistics** of S with respect to T can be computed in linear time using the generalized suffix array of S and T (i.e., the suffix array of $S\$T\$$) and its LCP array (exercise).

185

We will next describe the RMQ data structure for an arbitrary array $L[1..n]$ of integers.

- We precompute and store the minimum values for the following collection of ranges:
 - Divide $L[1..n]$ into blocks of size $\log n$.
 - For all $0 \leq \ell \leq \log(n/\log n)$, include all ranges that consist of 2^ℓ blocks. There are $\mathcal{O}(\log n \cdot \frac{n}{\log n}) = \mathcal{O}(n)$ such ranges.
 - Include all prefixes and suffixes of blocks. There are a total of $\mathcal{O}(n)$ of them.
- Now any range $L[i..j]$ that crosses or touches a block boundary can be exactly covered by at most **four ranges** in the collection.



The minimum value in $L[i..j]$ is the minimum of the minimums of the covering ranges and can be computed in constant time.

187

Enhanced Suffix Array

The enhanced suffix array adds two more arrays to the suffix and LCP arrays to make the data structure fully equivalent to suffix tree.

- The idea is to represent a suffix tree node v representing a factor S_v by the suffix array interval of the suffixes that begin with S_v . That interval contains exactly the suffixes that are in the subtree rooted at v .
- The additional arrays support navigation in the suffix tree using this representation: one array along the regular edges, the other along suffix links.

With all the additional arrays the suffix array is not very space efficient data structure any more. Nowadays suffix arrays and trees are often replaced with **compressed text indexes** that provide the same functionality in much smaller space.

189

Here are some of the key properties of the BWT.

- The BWT is easy to compute using the suffix array:

$$L[i] = \begin{cases} \$ & \text{if } SA[i] = 0 \\ T[SA[i] - 1] & \text{otherwise} \end{cases}$$

- The BWT is **invertible**, i.e., T can be reconstructed from the BWT L alone. The inverse BWT can be computed in the same time it takes to sort the characters.
- The BWT L is typically **easier to compress** than the text T . Many text compression algorithms are based on compressing the BWT.
- The BWT supports **backward searching**, a different technique for indexed exact string matching. This is used in many **compressed text indexes**.

191

Range Minimum Queries

The **range minimum query** (RMQ) asks for the smallest value in a given range in an array. Any array can be preprocessed in linear time so that RMQ for any range can be answered in constant time.

We can answer **longest common extension** (LCE) queries using RMQ queries on the LCP array:

Lemma 4.9: The length of the **longest common prefix** of two suffixes $T_i < T_j$ is $lcp(T_i, T_j) = \min\{LCP[k] \mid k \in [SA^{-1}[i] + 1..SA^{-1}[j]]\}$.

The lemma can be seen as a generalization of Lemma 1.31(b) (Lecture 3) and holds for any sorted array of strings. The proof is left as an exercise.

- In addition to the many general applications of LCE queries, we can also replace the LLCP and RLCP arrays in binary searching.

186

Ranges $L[i..j]$ that are completely inside one block are handled differently.

- Let $NSV(i) = \min\{k > i \mid L[k] < L[i]\}$ (NSV=Next Smaller Value). Then the position of the minimum value in the range $L[i..j]$ is the last position in the sequence $i, NSV(i), NSV(NSV(i)), \dots$ that is in the range. We call these the NSV positions for i .
- For each i , store the NSV positions for i up to the end of the block containing i as a bit vector $B(i)$. Each bit corresponds to a position within the block and is one if it is an NSV position. The size of $B(i)$ is $\log n$ bits and we can assume that it fits in a single machine word. Thus we need $\mathcal{O}(n)$ words to store $B(i)$ for all i .
- The position of the minimum in $L[i..j]$ is found as follows:
 - Turn all bits in $B(i)$ after position j into zeros. This can be done in constant time using bitwise shift-operations.
 - The right-most 1-bit indicates the position of the minimum. It can be found in constant time using a lookup table of size $\mathcal{O}(n)$.

All the data structures can be constructed in $\mathcal{O}(n)$ time (exercise).

188

Burrows–Wheeler Transform

The Burrows–Wheeler transform (BWT) is an important technique for **text compression**, **text indexing**, and their combination **compressed text indexing**.

Let $T[0..n]$ be the text with $T[n] = \$$. For any $i \in [0..n]$, $T[i..n]T[0..i]$ is a **rotation** of T . Let \mathcal{M} be the matrix, where the rows are all the rotations of T in lexicographical order. All columns of \mathcal{M} are **permutations** of T . In particular:

- The first column F contains the text characters in order.
- The last column L is the BWT of T .

Example 4.10: The BWT of $T = \text{banana}\$$ is $L = \text{annb}\$aa$.

F		L
\$	b a n a n a	a
a	\$ b a n a n	n
a	n a \$ b a n	n
a	n a n a \$ b	b
b	a n a n a \$	a
n	a \$ b a n a	a
n	a n a \$ b a	a

190

Inverse BWT

Let \mathcal{M}' be the matrix obtained by rotating \mathcal{M} one step to the right.

Example 4.11:

\mathcal{M}	\mathcal{M}'
\$ b a n a n a a	a \$ b a n a n a
a \$ b a n a n n	n a \$ b a n a n a
a n a \$ b a n n	n a n a \$ b a
a n a n a \$ b b	b a n a n a \$
b a n a n a \$ a	\$ b a n a n a
n a \$ b a n a a	a n a \$ b a n
n a n a \$ b a a	a n a n a \$ b

- The rows of \mathcal{M}' are the rotations of T in a different order.
- In \mathcal{M}' without the first column, the rows are sorted lexicographically. If we sort the rows of \mathcal{M}' **stably** by the first column, we obtain \mathcal{M} .

This cycle $\mathcal{M} \xrightarrow{\text{rotate}} \mathcal{M}' \xrightarrow{\text{sort}} \mathcal{M}$ is the key to inverse BWT.

192

Algorithm 4.15: Backward Search

Input: array C , function $rank_L$, pattern P

Output: suffix array range $[b..e)$ containing starting positions of P

```
(1)  $b \leftarrow 0$ ;  $e \leftarrow n + 1$ 
(2) for  $i \leftarrow m - 1$  downto 0 do
(3)    $c \leftarrow P[i]$ 
(4)    $b \leftarrow C[c] + rank_L(c, b)$ 
(5)    $e \leftarrow C[c] + rank_L(c, e)$ 
(6) return  $[b..e)$ 
```

- The array C requires an integer alphabet that is not too large.
- The trivial implementation of the function $rank_L$ as an array requires $\Theta(\sigma n)$ space, which is often too much. There are much more space efficient (but slower) implementations. There are even implementations with a size that is close to the size of the [compressed text](#). Such an implementation is the key component in many compressed text indexes. These are covered in the course [Data Compression Techniques](#).