

Tutorial 11

Data Structures 1, 9. 5. 2025

<https://iuuk.mff.cuni.cz/~chmel/2425/ds1en/>

Exercise 1 (Suffix and LCP arrays in practice)

Together, we will trace the building algorithms for both arrays for the string `barokoarokoko`. After that, you can try it yourselves for the strings `CAGTAGCTGTA`, `TATGTCAGTATCTC`.

Exercise 2 (Yes, this story is more than two years old)

You managed to save a bunch of elephants from a volcano¹. Thanks to your expert knowledge of various elephant species, you noticed that you actually have two different species of elephants. Since each species has different dietary preferences, you want to separate them so that none of the elephants get upset with you. Unfortunately, you are not able to distinguish them directly, but you know that they differ in their DNA. Moreover, as true computer science students, you have a DNA sequencing device with you, which also automatically generates the suffix and LCP arrays of the given DNA string.

With each elephant's consent, you sequence the relevant portion of its DNA, which can be represented as a string. After thoroughly studying the manual *How Elephants Differ*, you discovered that the main difference in their DNA is the length of the longest repeating substring of DNA bases. Specifically, if the length of such a substring is odd, the elephant is an Indian elephant, whereas an African elephant has its longest such substring of even length. For each elephant, determine whether it is an African or Indian elephant.

If you're not interested in the story: given a string and its suffix and LCP arrays, find the length of the longest repeating substring.

Exercise 3 (In search of palindromes)

Given a string α , find, using the LCP array or suffix array the longest palindromic substring of α .

Exercise 4 (The elephants are somewhat curious)

When the elephants found out everything your sequencer can do, they became interested in one more thing: they read that a substring P of length n is associated with increased intelligence, and so they are curious whether they have such a predisposition. Given the relevant portion of their genome G of length m , first recall the trivial algorithm for finding a needle in a haystack in time $\mathcal{O}(n \log m)$ (you may assume the suffix array is already constructed).

Next, suppose that for any two strings x, y , you are able to compute $\text{LCP}(x, y)$ in constant time. Modify the trivial algorithm so that it runs in only $\mathcal{O}(\log m)$ time.

Bonusy

Exercise 5 (Advent of code intensifies)

But the elephants are also realists, so they are aware that the general LCP cannot be computed in constant time. On the other hand, they have discovered a bunch of other substrings that are allegedly associated with all kinds of traits, so they would be very happy if the algorithm could be as fast as possible.

But what if we could compute the LCP of any two *suffixes* in the genome (haystack) in constant time?² In that case, find an algorithm for searching for a needle of length n in a haystack of length m that runs in time $\mathcal{O}(n + \log m)$.

Hint: In the suffix array, all suffixes of G are sorted lexicographically, so we can perform binary search on it, and we are interested in finding a suffix whose prefix is exactly P . The key is to ensure that, during comparisons, we only compare the relevant parts of the strings so that all comparisons together cost only $\mathcal{O}(n)$. The goal is to use information about common prefixes between elements of the suffix array to maintain information about the common prefix of P with relevant strings in S .

Exercise 6 (Actually, the assumption can be weakened)

We know that it is possible to compute the LCP array when given S in linear time (using Kasai's algorithm). Show, that we can then also compute all possible values that the algorithm in the previous exercise uses.

¹See the story from Advent of Code 2022, Days 16–18: <https://adventofcode.com/2022/day/16>.

²This assumption is stronger than just having the LCP array, because we are also able to compute the LCP of lexicographically non-adjacent suffixes.

Cheatsheet

- For a string α with Python-style substring notation (i.e., indexing starts at zero, $\alpha[i : j]$ includes all characters from index i to $j - 1$, omitting one coordinate means "from the beginning" or "to the end", respectively), we have the following:
- The suffix array S : $S[i]$ gives the index such that $\alpha[S[i] :]$ is the i -th lexicographically smallest suffix of the string α
- The rank array R : $R[i]$ gives the lexicographic rank of the suffix $\alpha[i :]$
- The LCP array L : $L[i]$ gives the number of characters at the beginning that are shared by $\alpha[S[i] :]$ and $\alpha[S[i + 1] :]$ (i.e., the i -th and $(i + 1)$ -th lexicographically smallest suffixes of α)