

Deterministic Sub-Linear Space LCE Data Structures With Efficient Construction*

Yuka Tanimura¹, Tomohiro I², Hideo Bannai³, Shunsuke Inenaga⁴, Simon J. Puglisi⁵, and Masayuki Takeda⁶

1 Department of Informatics, Kyushu University, Japan
yuka.tanimura@inf.kyushu-u.ac.jp

2 Kyushu Institute of Technology, Japan
tomohiro@ai.kyutech.ac.jp

3 Department of Informatics, Kyushu University, Japan
bannai@inf.kyushu-u.ac.jp

4 Department of Informatics, Kyushu University, Japan
inenaga@inf.kyushu-u.ac.jp

5 Department of Computer Science, University of Helsinki, Finland
puglisi@cs.helsinki.fi

6 Department of Informatics, Kyushu University, Japan
takeda@inf.kyushu-u.ac.jp

Abstract

Given a string S of n symbols, a longest common extension query $\text{LCE}(i, j)$ asks for the length of the longest common prefix of the i th and j th suffixes of S . LCE queries have several important applications in string processing, perhaps most notably to suffix sorting. Recently, Bille et al. (J. Discrete Algorithms 25:42–50, 2014, Proc. CPM 2015:65–76) described several data structures for answering LCE queries that offers a trade-off between data structure size and query time. In particular, for a parameter $1 \leq \tau \leq n$, their best deterministic solution is a data structure of size $O(\frac{n}{\tau})$ which allows LCE queries to be answered in $O(\tau)$ time. However, the construction time for all deterministic versions of their data structure is quadratic in n . In this paper, we propose a deterministic solution that achieves a similar space-time trade-off of $O(\tau \min\{\log \tau, \log \frac{n}{\tau}\})$ query time using $O(\frac{n}{\tau})$ space, but we significantly improve the construction time to $O(n\tau)$.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases longest common extension, longest common prefix, sparse suffix array

Digital Object Identifier 10.4230/LIPIcs.CPM.2016.1

1 Introduction

Given a string S of n symbols, a *longest common extension query* $\text{LCE}(i, j)$ asks for the length of the longest common prefix of the i th and j th suffixes of S .

The ability to efficiently answer LCE queries allows optimal solutions to many string processing problems. Gusfield’s book [4], for example, lists several applications of LCEs to basic pattern matching and discovery problems, including: pattern matching with wildcards, mismatches and errors; the detection of various types of palindromes (maximal, complimented, separated, approximate); and the detection of repetitions and approximate repetitions.

* HB, SI, MT were supported by JSPS KAKENHI Grant Numbers 25280086, 26280003, 25240003.



© Yuka Tanimura, Tomohiro I, Hideo Bannai, Shunsuke Inenaga, Simon J. Puglisi, and Masayuki Takeda;
licensed under Creative Commons License CC-BY

27th Annual Symposium on Combinatorial Pattern Matching (CPM 2016).

Editors: Roberto Grossi and Moshe Lewenstein; Article No. 1; pp. 1:1–1:10



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

■ **Table 1** Deterministic solutions to LCE.

Data Structure		Preprocessing		Trade-off range	Reference
Space	Query	Space	Time		
1	n	1	1	-	naïve computation
n	1	n	n	-	suffix array + RMQ
$\frac{n}{\tau}$	τ^2	$\frac{n}{\tau}$	$\frac{n^2}{\tau}$	$1 \leq \tau \leq \sqrt{n}$	[3]
$\frac{n}{\tau}$	$\tau \log^2 \frac{n}{\tau}$	$\frac{n}{\tau}$	n^2	$1 \leq \tau \leq n$	[2], Section 2
$\frac{n}{\tau}$	τ	$\frac{n}{\tau}$	$n^{2+\epsilon}$	$1 \leq \tau \leq n$	[2], Section 4
$\frac{n}{\tau}$	$\tau \log^2 \frac{n}{\tau}$	$\frac{n}{\tau}$	$n\tau + n \log \frac{n}{\tau}$	$1 \leq \tau \leq n$	This work, Theorem 9
$\frac{n}{\tau}$	$\tau \log \tau$	$\frac{n}{\tau}$	$n\tau$	$1 \leq \tau \leq \frac{n}{\log n}$	This work, Theorem 10
$\frac{n}{\tau}$	$\tau \min\{\log \tau, \log \frac{n}{\tau}\}$	$\frac{n}{\tau}$	$n\tau$	$1 \leq \tau \leq n$	This work, Corollary 12

Lempel-Ziv parsing [6] and suffix sorting [7, 5] are two more fundamental string processing problems to which LCEs are key.

Without preprocessing, answering an arbitrary query $\text{LCE}(i, j)$ requires $O(n)$ time: we simply compare the suffixes starting at positions i and j character by character until we find a mismatch. To answer queries faster we could build the suffix tree and preprocess it for lowest-common-ancestor queries. This well-known solution answers queries in $O(1)$ time and the data structure is of $O(n)$ size and takes $O(n)$ time to construct.

In recent years, motivated by scenarios where $O(n)$ space is prohibitive, several authors have sought data structures that achieve a trade-off between data structure size and query time [13, 3, 2]. The best trade-off to date is due to Bille et al. [2], where they describe a data structure of size $O(n/\tau)$ which allow LCE queries to be answered in $O(\tau)$ time.

However, as described in [2], their deterministic data structure requires $O(n^2)$ time to construct if only $O(n/\tau)$ working space is allowed. This is a major drawback, because it does not allow the space-query time trade-off to be passed on to applications—indeed, construction of the data structure would become a time bottleneck in all the applications listed above. We note that Bille et al. [2] also proposed randomized solutions which achieve the same space-query time trade-off with subquadratic preprocessing time. In this paper, we focus on deterministic solutions.

The main contributions of this article are as follows:

1. We describe a new data structure for LCEs that has size $O(\frac{n}{\tau})$, query time $O(\tau \log \tau)$, and, critically, can be constructed in $O(n\tau)$ time.
2. We show how to combine the new data structure with one of Bille et al. to derive a structure that has $O(\tau \min\{\log \tau, \log \frac{n}{\tau}\})$ query time and the same space and construction bounds as the new structure. As a side result, we also show how this particular structure of Bille et al. can be constructed efficiently.

Table 1 summarizes our results and previous work on the deterministic version of the problem.

In the next section we lay down notation and some basic algorithmic and data structural tools. Then, in Section 3, we introduce our new LCE data structures, beginning with a slightly modified version of one of Bille et al.’s data structures, followed by the new and combined data structures. Section 4 deals with efficient construction. We finish, in Section 5, by noting that our new structures lead directly to improved (deterministic) bounds for the sparse suffix sorting problem.

2 Preliminaries

Let $\Sigma = \{1, \dots, \sigma\}$ denote the alphabet, and Σ^* the set of strings. If $w = xyz$ for some strings w, x, y, z , then x, y , and z are respectively called a *prefix*, *substring*, and *suffix* of w . For any string w , let $|w|$ denote the length of w , and for any $0 \leq i < |w|$, let $w[i]$ denote the i th character of w , i.e., $w = w[0] \cdots w[|w| - 1]$. For convenience, let $w[i] = 0$ when $i \geq |w|$. For any $0 \leq i \leq j$, let $w[i..j] = w[i] \cdots w[j]$, and for any $0 \leq i < |w|$, let $w[i..] = w[i..|w| - 1]$. We denote $x \prec y$ if a string x is lexicographically smaller than a string y .

For any string w , let $\text{lcp}_w(i, j)$ denote the length of the longest common prefix of $w[i..]$ and $w[j..]$. We will write $\text{lcp}(i, j)$ when w is clear from the context. Since $\text{lcp}_w(i, i) = |w| - i$, we will only consider the case when $i \neq j$. Note that answering an LCE query $\text{LCE}(i, j)$ is equivalent to computing $\text{lcp}_w(i, j)$.

For any integers $i \leq j$, let $[i..j]$ denote the set of integers from i to j (including i and j), and for $0 \leq p < \tau$, let $[i..j]_p^\tau = \{k \mid k \in [i..j], k \bmod \tau = p\}$.

For any string w of length n and $0 \leq p < \tau$, let $\hat{w}_{\tau,p}$ denote a string of length $\lceil (|w| - p) / \tau \rceil$ over the alphabet $\{1, \dots, \sigma^\tau\}$ such that $\hat{w}_{\tau,p}[i] = w[p + \tau i..p + \tau(i + 1) - 1]$ for any $i \geq 0$. We call $\hat{w}_{\tau,p}$ the *meta-string* of w wrt. sampling rate τ and offset p , and each character of $\hat{w}_{\tau,p}$ is called a *meta-character*.

In the rest of the paper, we assume a polynomially bounded integer alphabet, i.e., for some constant $c \geq 0$, $\sigma = O(n^c)$ for any input string w of length n .

► **Definition 1** ([12]). The suffix array SA_w of a string w of length n is an array of size n containing a permutation of $[0..n - 1]$ that represents the lexicographic order of the suffixes of w , i.e., $w[\text{SA}_w[0]..] \prec \cdots \prec w[\text{SA}_w[n - 1]..]$. The inverse suffix array ISA_w is an array of size n such that $\text{ISA}_w[\text{SA}_w[i]] = i$ for all $0 \leq i < n$. The LCP array LCP_w of a string w of length n is an array of size n such that $\text{LCP}_w[0] = 0$ and $\text{LCP}_w[i] = \text{lcp}_w(\text{SA}_w[i - 1], \text{SA}_w[i])$ for $0 < i < n$.

► **Lemma 2** ([9, 10, 11, 7]). For any string w of length n , the arrays $\text{SA}_w, \text{ISA}_w, \text{LCP}_w$ can be computed in $O(n)$ time and space.

For any array A and $0 \leq i \leq j < |A|$, let $\text{rmq}_A(i, j)$ denote a Range Minimum Query (RMQ), i.e., $\text{rmq}_A(i, j) = \arg \min_{k \in [i..j]} \{A[k]\}$. It is well known that A can be preprocessed in linear time and space so that $\text{rmq}_A(i, j)$, for any $0 \leq i \leq j < |A|$, can be answered in constant time [1]. Since $\text{lcp}_w(i, j) = \text{LCP}_w[\text{rmq}_{\text{LCP}_w}(i' + 1, j')]$ where $i' = \min\{\text{ISA}_w(i), \text{ISA}_w(j)\}$ and $j' = \max\{\text{ISA}_w(i), \text{ISA}_w(j)\}$, it follows that a string of length n can be preprocessed in $O(n)$ time and space so that for any $0 \leq i, j < n$, $\text{lcp}_w(i, j)$ can be computed in $O(1)$ time.

Our algorithm relies on sparse suffix arrays. For a string w of length n and any set $P \subseteq [0..n - 1]$ of positions, let $\text{SSA}_P[0..|P| - 1]$ be an array consisting of entries of SA that are in P , i.e., for any $0 \leq i < |P|$, $\text{SSA}_P[i] \in P$, and $w[\text{SSA}_P[0]..] \prec \cdots \prec w[\text{SSA}_P[|P| - 1]..]$. The sparse LCP array $\text{SLCP}_P[0..|P| - 1]$ is defined analogously, $\text{SLCP}_P[i] = \text{lcp}_w(\text{SSA}_P[i - 1], \text{SSA}_P[i])$.

Let $1 \leq \tau \leq n$ be a parameter called the *sampling rate*. When, $P = [0..n - 1]_p^\tau$, for some $0 \leq p < \tau \leq n$, SSA_P is called the evenly spaced sparse suffix array with sampling rate τ and offset p . Given an evenly spaced sparse suffix array SSA_P , we can compute in $O(\frac{n}{\tau})$ time, a representation of the sparse inverse suffix array ISA_P as an array X of size $O(\frac{n}{\tau})$ where $X[\lfloor \text{SSA}_P[i] / \tau \rfloor] = i$, i.e., $\text{ISA}_P[i] = X[\lfloor i / \tau \rfloor]$ for all $i \in P$. By directly applying the algorithm of Kasai et al. [9], SLCP_P can be computed from SSA_P and (the representation of) ISA_P in $O(n)$ time and $O(\frac{n}{\tau})$ space.

3 Data structure and query computation

Our algorithms are based on the same observation as used in [2].

► **Observation 3** ([2]). *For any positions $i, j, k \in [0..n - 1]$ if $\text{lcp}(j, k) \geq \text{lcp}(i, j)$ then, $\text{lcp}(i, j) = \min\{\text{lcp}(i, k), \text{lcp}(j, k)\}$.*

The observation allows us to reduce the computation of lcp values between a pair of positions, to the computation of lcp values between another pair of values, both from a specific subset of positions. For each specific position i , called a sampled position, and for each such subset S , a position $\pi(i, S) = \arg \max_{i' \in S} \{\text{lcp}(i, i')\}$ is precomputed. The idea is that the size of S gets smaller after each reduction, therefore giving a bound on the query time.

► **Corollary 4.** *For any pair of positions $i \in S \subseteq [0..n - 1]$ and $j \in [0..n - 1]$, $\text{lcp}(i, j) = \min\{\text{lcp}(i, \pi(j, S)), \text{lcp}(j, \pi(j, S))\}$.*

3.1 Bille et al.'s data structure

We first introduce a slightly modified version of the deterministic data structure by Bille et al. [2] that uses $O(\frac{n}{\tau})$ space and allows queries in $O(\tau \log^2 \frac{n}{\tau})$ time, where τ is a parameter in the range $1 \leq \tau \leq n$. We note that the modifications do not affect the asymptotic complexities.


Let $t = \tau \lceil \log \frac{n}{\tau} \rceil$, $p = (n - 1) \bmod t$ and let $\mathcal{P} = [0..n - 1]_p^t$ be the set of positions called *sampled positions*. The data structure of [2] to compute $\text{lcp}(i, j)$ for any $0 \leq i < j < n$ consists of two main parts, one for when $j - i \geq t$, and the other for when $j - i < t$. Since we will use the latter part as is, we will only describe the former. The query time, space, and preprocessing time of the latter part are respectively, $O(\tau \log \frac{n}{\tau})$, $O(\frac{n}{\tau})$, and $O(n)$ (see Section 2 of [2]).

Consider a full binary tree where the root corresponds to the interval $[0..n - 1]$, and for any node, the left and right children split their parent interval almost evenly, but assuring that the right-most position in the left child is a sampled position. Thus, there will be $\lceil n/t \rceil$ leaves corresponding to intervals of size t (except perhaps for the leftmost interval which may be smaller), and the height of the tree is $O(\log \frac{n}{t})$. For any internal node v in the tree, let I_v denote its corresponding interval, and $\ell(v), r(v)$ respectively the left and right children of v . For all sampled positions $i \in I_{r(v)} \cap \mathcal{P}$, a position $\pi(i, I_{\ell(v)}) = \arg \max_{i' \in I_{\ell(v)}} \{\text{lcp}(i, i')\}$ and a value $L(i, I_{\ell(v)}) = \text{lcp}(i, \pi(i, I_{\ell(v)}))$ are computed and stored. The size of the data structure is therefore $O(\frac{n}{t} \log \frac{n}{t}) = O(\frac{n}{\tau})$.

Assume w.l.o.g. that $j > i$. A query for $\text{lcp}(i, j)$ with $j - i \geq t$ is computed as follows. First, compare up to $\delta < t$ characters of $w[i..]$ and $w[j..]$ until we encounter a mismatch, in which case we obtain an answer, or $j + \delta$ is a sampled position. Let I_v be the interval such that $i + \delta \in I_{\ell(v)}$ and $j + \delta \in I_{r(v)}$. From the preprocessing, we obtain a position $\pi(j + \delta, I_{\ell(v)}) \in I_{\ell(v)}$, which, from Corollary 4, satisfies:

$$\begin{aligned} \text{lcp}(i, j) &= \delta + \text{lcp}(i + \delta, j + \delta) \\ &= \delta + \min\{\text{lcp}(i + \delta, \pi(j + \delta, I_{\ell(v)})), \text{lcp}(j + \delta, \pi(j + \delta, I_{\ell(v)}))\} \\ &= \delta + \min\{\text{lcp}(i + \delta, \pi(j + \delta, I_{\ell(v)})), L(j + \delta, I_{\ell(v)})\} \end{aligned}$$

Thus, the problem can be reduced to computing $\text{lcp}(i + \delta, \pi(j + \delta, I_{\ell(v)}))$, where both $i + \delta, \pi(j + \delta, I_{\ell(v)}) \in I_{\ell(v)}$, and we apply the algorithm recursively. Note that if $j \in I_{r(v)}$ we have, from the definition of the intervals, that $j + \delta \in I_{r(v)}$, so each recursion takes us



■ **Figure 1** Sets S_k , with $k = 1, 2, 3$, for the sampled positions specified by black dots.

further down the tree. When an interval corresponding to a leaf node is reached, we have that $j - i < t$ and use the other data structure (for a description of which we refer the reader to [2]). Since we compare up to t characters at each level, the total query time is $O(t \log \frac{n}{t}) = O(\tau \log^2 \frac{n}{\tau})$.

3.2 New data structure

Let $t = \tau \lceil \log \tau \rceil$, $p = (n - 1) \bmod t$, and let $\mathcal{P} = [0..n - 1]_p^t$ be the set of sampled positions. Instead of considering a hierarchy of intervals of positions, we classify the positions according to their distance to the closest sampled position to their right. Define $S_k = \{i \mid (i + d) \bmod t = p, d \in ([2^{k-1}..2^k - 1] \cap [1..t - 1])\}$ for $k = 1, \dots, \lceil \log t \rceil$ (see also Figure 1).

The preprocessing computes and stores for each sampled position $i \in \mathcal{P}$ and each S_k , a position $\pi(i, S_k) = \arg \max_{i' \in S_k} \{\text{lcp}(i, i')\}$, and a value $L(i, S_k) = \text{lcp}(i, \pi(i, S_k))$. Also, $\text{SLCP}_{\mathcal{P}}$ is computed and preprocessed for range minimum queries so that for any $i, j \in \mathcal{P}$, $\text{lcp}(i, j)$ can be computed in constant time. Thus, the space required for the data structure is $O(\frac{n}{t} \log t) = O(\frac{n}{\tau})$.

A value $\text{lcp}(i, j)$ is computed as follows. First, compare up to δ characters of $w[i..]$ and $w[j..]$ until we encounter a mismatch, in which case we obtain an answer, or, either $i + \delta$ or $j + \delta$ is a sampled position. If both $i + \delta$ and $j + \delta$ are sampled positions, $\text{lcp}(i, j) = \delta + \text{lcp}(i + \delta, j + \delta)$ can be answered in constant time. Assume w.l.o.g. that only $j + \delta$ is a sampled position, and let k be such that $i + \delta \in S_k$. Then, from Corollary 4 and the preprocessing, we have

$$\begin{aligned} \text{lcp}(i, j) &= \delta + \text{lcp}(i + \delta, j + \delta) \\ &= \delta + \min\{\text{lcp}(i + \delta, \pi(j + \delta, S_k)), \text{lcp}(j + \delta, \pi(j + \delta, S_k))\} \\ &= \delta + \min\{\text{lcp}(i + \delta, \pi(j + \delta, S_k)), L(j + \delta, S_k)\} \end{aligned}$$

and the problem has been reduced to computing $\text{lcp}(i + \delta, \pi(j + \delta, S_k))$ where both $i + \delta, \pi(j + \delta, S_k) \in S_k$, and the processes are repeated. Notice that in the next step, at least 2^{k-1} characters are compared until one of the two positions becomes a sampled position. This implies that the remaining distance to the closest sampled position of the other position will be at most $2^{k-1} - 1$, and thus the position will be in $S_{k'}$ for some $k' \leq k - 1$. Therefore, the process will only be repeated at most $\lceil \log t \rceil$ times. Because the number of characters compared in each step is bounded by t and is at least halved every step, the total number of character comparisons, and thus the query time, is $O(t) = O(\tau \log \tau)$.

3.3 Combining the structures

We can combine the structures described in Sections 3.1 and 3.2, to achieve $O(\tau \log \frac{n}{\tau})$ query time using $O(\frac{n}{\tau})$ space for $1 \leq \tau \leq n$. Furthermore, we can achieve $O(\tau \min\{\log \tau, \log \frac{n}{\tau}\})$ query time by choosing the better structure depending on τ . More precisely, when $\tau \leq \frac{n}{\tau}$ (i.e., $\tau \leq \sqrt{n}$), we simply use the structure of Section 3.2, and when $\tau \geq \frac{n}{\tau}$ (i.e., $\tau \geq \sqrt{n}$), we use the combined structure. Thus, we assume below that $\tau \geq \frac{n}{\tau}$.

Let $t = \tau \lceil \log \frac{n}{\tau} \rceil$, $p = (n - 1) \bmod t$ and let $\mathcal{P} = [0..n - 1]_p^t$ be the set of positions called *sampled positions*. We consider both the structures described in Section 3.1 and Section 3.2, with the following modifications. Let $d_t = 2^{\lceil \log t \rceil - \lceil \log \frac{n}{t} \rceil} = O(\frac{t^2}{n})$. For Bille et al.'s data structure, we make two modifications. First, for each node I_v and sampled position $i \in I_{r(v)} \cap \mathcal{P}$, we only consider points that are at most d_t from the closest sampled position to the right, i.e., instead of $\pi(i, I_{\ell(v)})$ and $L(i, I_{\ell(v)})$, we compute and store a position $\pi(i, I_{\ell(v)} \cap D) = \arg \max_{i' \in I_{\ell(v)} \cap D} \{\text{lcp}(i, i')\}$ and a value $L(i, I_{\ell(v)} \cap D)$, where $D = \{i' \mid (i' + d) \bmod t = p, 0 \leq d < d_t\}$. In addition to this, we compute and store for all sampled position $i \in I_{\ell(v)} \cap \mathcal{P}$, a position $\pi(i, I_{r(v)} \cap D) = \arg \max_{i' \in I_{r(v)} \cap D} \{\text{lcp}(i, i')\}$ and $L(i, I_{r(v)} \cap D) = \text{lcp}(i, \pi(i, I_{r(v)} \cap D))$. This will only double the total size of the structure and thus the space usage remains $O(\frac{n}{\tau})$. For the new data structure, we keep the definitions of $\pi(i, S_k)$ and $L(i, S_k)$, but store these values only for $k = \lceil \log t \rceil - \lceil \log \frac{n}{t} \rceil, \dots, \lceil \log t \rceil$. Thus, although the value of t has changed, the total size of the data structure is still $O(\frac{n}{t} \log \frac{n}{t}) = O(\frac{n}{\tau})$.

Queries $\text{lcp}(i, j)$ are answered as follows: First use the new data structure recursively using the original algorithm until the problem is reduced to a query between a sampled position and another position not in any S_k ($k \in [\lceil \log t \rceil - \lceil \log \frac{n}{t} \rceil .. \lceil \log t \rceil]$). This means that the distance from either of the query positions to the closest sampled position is at most d_t . The total number of character comparisons conducted is $O(t) = O(\tau \log \frac{n}{\tau})$. Then, we switch to Bille et al.'s structure using the original algorithm with the exception that we continue until *either* $i + \delta$ *or* $j + \delta$ (instead of just $j + \delta$) is a sampled position when comparing up to δ characters of $w[i..]$ and $w[j..]$. Since the distance to the closest sampling position is at most $O(\frac{t^2}{n})$ and by definition of $\pi(i, I_{\ell(v)} \cap D)$ and $\pi(i, I_{r(v)} \cap D)$, we have that this condition holds for all following recursive calls. Thus, at most $O(\frac{t^2}{n})$ character comparisons will be conducted at each level, for a total of $O(\frac{t^2}{n} \log \frac{n}{t}) = O(t(\frac{n}{t})^{-1} \log \frac{n}{t}) = O(t) = O(\tau \log \frac{n}{\tau})$.

4 Building the structures

Bille et al. [2] describe a preprocessing that runs in $O(n^2)$ time¹ and $O(\frac{n}{\tau})$ space. Here, we show that this can be reduced to $O(\tau n + n \log \frac{n}{\tau})$ time using the same space. While the algorithm of [2] builds the sparse suffix array containing only the suffixes starting at sampled positions and applies pattern matching, our trick is to build a sparse suffix array and sparse LCP array that includes other suffixes as well, in several (namely τ) rounds, so that the suffixes with maximum LCP with respect to each sampled position can be found by scans of the suffix array.

For integer alphabets, sparse suffix arrays and sparse LCP arrays can be constructed in $O(n)$ time if $O(n)$ space is allowed, simply by first building the (normal) suffix array and LCP array and removing the unwanted elements. For constant size alphabets, the evenly spaced sparse suffix array and sparse LCP array with sampling rate τ can be constructed in $O(n)$ time and $O(\frac{n}{\tau})$ space [8]. However, when the alphabet size σ is not constant, this is $O(n \log \sigma)$ time and $O(\frac{n}{\tau})$ space, since the computation is based on character comparisons. (Notice that simple application of linear time algorithms for computing the suffix array for

¹ However, we believe the analysis in Section 2.5 of [2] is not entirely correct; although the size of $|I|$ is halved at each level, their numbers double, and so the time complexity should be $O(n \cdot n + n \cdot (n/2) \cdot 2 \cdots + n \cdot (n/t) \cdot t) = O(n^2 \log \frac{n}{t})$ time. Also, they assume that the evenly spaced sparse suffix array can be constructed in $O(n)$ time and $O(\frac{n}{\tau})$ space for the integer alphabet. However, the paper they cite assumes a constant size alphabet and to the best of our knowledge, we do not know of an algorithm achieving such space-time trade-off.

the meta string will not achieve $O(n)$ time and $O(\frac{n}{\tau})$ space, since the use of radix sort implies $\Omega(\sigma)$ space for the buckets.) Repeated τ times, this results in $O(n\tau \log \sigma)$ time using $O(\frac{n}{\tau})$ space.

We first describe a technique to compute the sparse suffix array and the corresponding LCP array that contains two sets of evenly spaced suffixes, namely for offsets p and q , and to repeat this τ times, namely for offsets $p = (n-1) \bmod \tau$ and $q = (n-1) \bmod \tau, \dots, (n-\tau) \bmod \tau$, so that the total time for their construction is $O(n\tau)$ time using $O(\frac{n}{\tau})$ space. Then, we describe the construction of the data structures of Section 3 using this technique.

4.1 Common tools

For any string (or meta-string) w and $0 \leq i < |w|$, let CA_w denote an array containing a permutation of $[0..|w| - 1]$ such that $w[\text{CA}_w[i]] \leq w[\text{CA}_w[j]]$ for any $0 \leq i < j < |w|$, i.e., CA_w is an array of positions sorted according to the character at each position. (Note that CA_w is not necessarily unique.)

► **Lemma 5.** *For any string w and $0 \leq p < \tau$, $\text{CA}_{\hat{w}_{\tau,p}}$ can be computed in $O(n \log \tau)$ time using $O(\frac{n}{\tau})$ space.*

Proof. Since each character of w can be represented in $O(\log n)$ bits, the length of each meta-character of $\hat{w}_{\tau,p}$ is $O(\tau \log n)$ bits. We simply use LSD radix sort with a bucket size of $\frac{n}{\tau}$, i.e., we bucket sort using $\log(n/\tau)$ bits at a time. Thus, $O(\frac{\tau \log n}{\log(n/\tau)})$ rounds of bucket sort is conducted on $\frac{n}{\tau}$ items, resulting in $O(\frac{n \log n}{\log(n/\tau)}) = O(\frac{n(\log \tau + \log(n/\tau))}{\log(n/\tau)}) = O(n \log \tau)$ time giving the result. ◀

► **Lemma 6.** *For any string w and $0 \leq p < \tau$, $\text{CA}_{\hat{w}_{\tau,p}}$ can be computed from $\text{CA}_{\hat{w}_{\tau,p'}}$, where $p' = (p+1) \bmod \tau$, in $O(\frac{n}{\tau} \log \tau)$ time and $O(\frac{n}{\tau})$ space.*

Proof. We simply continue the LSD radix sort, and do an extra $O(\frac{\log n}{\log(n/\tau)})$ rounds of bucket sort for the preceding character of each meta-character, which results in $O(\frac{n}{\tau} \cdot \frac{\log n}{\log(n/\tau)}) = O(\frac{n}{\tau} \cdot \frac{\log \tau + \log(n/\tau)}{\log(n/\tau)}) = O(\frac{n}{\tau} \log \tau)$ time. ◀

► **Lemma 7.** *For any string w , $0 \leq p, q < \tau$, let $P = [0..n-1]_p^\tau$ and $Q = [0..n-1]_q^\tau$. Given $\text{CA}_{\hat{w}_{\tau,p}}$ and $\text{CA}_{\hat{w}_{\tau,q}}$, $\text{SSA}_{P \cup Q}$ and $\text{SLCP}_{P \cup Q}$ can be computed in $O(n)$ time using $O(\frac{n}{\tau})$ space.*

Proof. We first compute $\text{CA}_{w'}$ for meta-string $w' = \hat{w}_{\tau,p}0\hat{w}_{\tau,q}$. This can be done in $O(n)$ time and $O(\frac{n}{\tau})$ space by merging $\text{CA}_{\hat{w}_{\tau,p}}$ and $\text{CA}_{\hat{w}_{\tau,q}}$, (and adding $|\hat{w}_{\tau,p}0|$ to entries in $\text{CA}_{\hat{w}_{\tau,q}}$) since each comparison of meta characters can be done in $O(\tau)$ time. Using $\text{CA}_{w'}$, we then rename the characters of w' and create a string w^* such that $w^*[i] = |\{w'[j] \mid w'[j] < w'[i], 0 \leq j < |w'|\}| + 1$, in $O(n)$ time and $O(\frac{n}{\tau})$ space. Since w^* consists of integers bounded by its length, we can apply any linear-time suffix sorting algorithm and compute SA_{w^*} and LCP_{w^*} in $O(\frac{n}{\tau})$ time and space. As the lexicographic order of suffixes of w^* (except for $\text{SSA}_{w^*}[0] = |\hat{w}_{\tau,p}|$) corresponds to the lexicographic order of suffixes of w that start at positions in $P \cup Q$, we can obtain $\text{SSA}_{P \cup Q}$ from SA_{w^*} by appropriately translating the indices. More precisely, for $1 \leq i < |w'|$, let $\text{SSA}_{w^*}[i] = j$. If $0 \leq j < |\hat{w}_{\tau,p}|$, then $\text{SSA}_{P \cup Q}[i-1] = j\tau + p$, and otherwise (if $|\hat{w}_{\tau,p}0| \leq j < |w'|$), then $\text{SSA}_{P \cup Q}[i-1] = (j - |\hat{w}_{\tau,p}0|)\tau + q$. We can also obtain $\text{SLCP}_{P \cup Q}$ from LCP_{w^*} by multiplying a factor of τ and doing up to τ character comparisons per pair of adjacent suffixes in the suffix array, in a total of $O(n)$ time. ◀

► **Corollary 8.** *For any string w , let $p = n \bmod \tau$. The arrays $\text{SSA}_{P \cup Q}$ and $\text{SLCP}_{P \cup Q}$ can be computed successively for each $q = p, (p-1) \bmod \tau, \dots, (p-\tau+1) \bmod \tau$, where $P = [0..n-1]_p^\tau$ and $Q = [0..n-1]_q^\tau$, in $O(n\tau)$ time using $O(\frac{n}{\tau})$ space.*

Proof. For $p = q$, we first compute $\text{CA}_{\hat{w}_{\tau,p}} = \text{CA}_{\hat{w}_{\tau,q}}$ using Lemma 5. By applying Lemma 6, we can successively compute $\text{CA}_{\hat{w}_{\tau,q}}$ for $q = (p-1) \bmod \tau, \dots, (p-\tau+1) \bmod \tau$. Thus, with Lemma 7, we can successively compute $\text{SSA}_{P \cup Q}$ and $\text{SLCP}_{P \cup Q}$ in $O(n\tau)$ total time and $O(\frac{n}{\tau})$ space. \blacktriangleleft

4.2 Faster construction of Bille et al.'s data structure

We show that Bille et al.'s data structure can be constructed in $O(n\tau + n \log \frac{n}{\tau})$ time using $O(\frac{n}{\tau})$ space. Let $p = (n-1) \bmod \tau$. Using Corollary 8, we successively compute $\text{SSA}_{P \cup Q}$ and $\text{SLCP}_{P \cup Q}$ for each $q = p, (p-1) \bmod \tau, \dots, (p-\tau+1) \bmod \tau$, where $P = [0..n-1]_p^\tau$ and $Q = [0..n-1]_q^\tau$. This can be done in a total of $O(n\tau)$ time, and $O(\frac{n}{\tau})$ space. Recall that $t = \tau \lceil \log \frac{n}{\tau} \rceil$, and $\mathcal{P} = [0..n-1]_{p'}^t$, where $p' = (n-1) \bmod t$. Since t is a multiple of τ , we have $\mathcal{P} \subseteq P$.

For each q we do the following: $\text{SLCP}_{P \cup Q}$ is preprocessed in $O(\frac{n}{\tau})$ time and space to answer RMQ in constant time, thus allowing us to compute $\text{lcp}(i, j)$ for any $i, j \in P \cup Q$ in constant time. For any interval $I_v \subseteq [0..n-1]$ corresponding to a node in the binary tree let $I_v^q = I_v \cap (P \cup Q)$. Note that for $I_{\text{root}} = [0..n-1]$, $\text{SSA}_{I_{\text{root}}^q} = \text{SSA}_{P \cup Q}$. Now, for any node I_v , assume that $\text{SSA}_{I_v^q}$ is already computed. By simple linear time scans on $\text{SSA}_{I_v^q}$, we can obtain, for each sampled position $i = \text{SSA}_{I_v^q}[x] \in I_{r(v)}^q \cap \mathcal{P}$, the two suffixes $\text{SSA}_{I_v^q}[j^-], \text{SSA}_{I_v^q}[j^+] \in I_{\ell(v)}^q \cap Q$ which are lexicographically closest to i , i.e., $j^- = \max\{j < x \mid \text{SSA}_{I_v^q}[j] \in I_{\ell(v)}^q \cap Q\}$, $j^+ = \min\{j > x \mid \text{SSA}_{I_v^q}[j] \in I_{\ell(v)}^q \cap Q\}$, if they exist. Then, the larger of $\text{lcp}(i, \text{SSA}_{I_v^q}[j^-])$ and $\text{lcp}(i, \text{SSA}_{I_v^q}[j^+])$ gives $\pi(i, I_{\ell(v)}^q \cap Q) = \arg \max_{i' \in I_{\ell(v)}^q \cap Q} \{\text{lcp}(i, i')\}$ and $L(i, I_{\ell(v)}^q \cap Q) = \text{lcp}(i, \pi(i, I_{\ell(v)}^q \cap Q))$. Since $i, \text{SSA}_{I_v^q}[j^+], \text{SSA}_{I_v^q}[j^-] \in P \cup Q$, these values can be computed in constant time, which is $O(|I_v^q|)$ total time for all sampled positions $i \in I_{r(v)}^q \cap \mathcal{P}$. Next, for the child intervals, $\text{SSA}_{I_{\ell(v)}^q}$ and $\text{SSA}_{I_{r(v)}^q}$ can be computed in $O(|I_v^q|)$ time by a simple scan on $\text{SSA}_{I_v^q}$, and the computation is performed recursively for each child. Since the union of $I_v^q \cap Q$ over all q is I_v , we have $\pi(i, I_{\ell(v)}) = \pi(i, I_{\ell(v)}^q)$ and $L(i, I_{\ell(v)}) = L(i, I_{\ell(v)}^q)$, where $\hat{q} = \arg \max_{0 \leq q' < \tau} \{\text{lcp}(i, \pi(i, I_{\ell(v)}^{q'} \cap Q))\}$, so we can obtain $\pi(i, I_{\ell(v)})$ and $L(i, I_{\ell(v)})$ for each sampled position i and interval I_v by repeating the above process for each q .

Since the processing at each node is linear in the size of the arrays whose total size at a given level is $O(\frac{n}{\tau})$, the total time for the recursion is $O(\frac{n}{\tau} \log \frac{n}{\tau})$ for each q . Thus in total, the preprocessing can be done in $O(n\tau + n \log \frac{n}{\tau})$ time.

► **Theorem 9.** *For any string of length n and integer $1 \leq \tau \leq n$, a data structure of size $O(\frac{n}{\tau})$ can be constructed in $O(n\tau + n \log \frac{n}{\tau})$ time using $O(\frac{n}{\tau})$ space, such that for any $0 \leq i, j < n$, $\text{lcp}(i, j)$ can be answered in $O(\tau \log^2 \frac{n}{\tau})$ time.*

4.3 Fast construction of new data structure

Let $p = (n-1) \bmod \tau$. Using Corollary 8, we successively compute $\text{SSA}_{P \cup Q}$ and $\text{SLCP}_{P \cup Q}$ for each $q = p, (p-1) \bmod \tau, \dots, (p-\tau+1) \bmod \tau$, where $P = [0..n-1]_p^\tau$ and $Q = [0..n-1]_q^\tau$. This can be done in a total of $O(n\tau)$ time, and $O(\frac{n}{\tau})$ space. Recall that $t = \tau \lceil \log \tau \rceil$, and $\mathcal{P} = [0..n-1]_{p'}^t$, where $p' = (n-1) \bmod t$. Since t is a multiple of τ , we have $\mathcal{P} \subseteq P$.

For each q we do the following: $\text{SLCP}_{P \cup Q}$ is preprocessed in $O(\frac{n}{\tau})$ time and space to answer RMQ in constant time, thus allowing us to compute $\text{lcp}(i, j)$ for $i, j \in P \cup Q$ in constant time. Let $S_k^q = S_k \cap Q$ for any $1 \leq k \leq \lceil \log t \rceil$. Next, we conduct for each $k = 1, \dots, \lceil \log t \rceil$, linear time scans on $\text{SSA}_{P \cup Q}$ so that for each sampled position $i = \text{SSA}_{P \cup Q}[x] \in \mathcal{P}$, the two suffixes $\text{SSA}_{P \cup Q}[j^-], \text{SSA}_{P \cup Q}[j^+] \in S_k^q$ which are lexicographically closest to i , i.e., $j^- = \max\{j < x \mid \text{SSA}_{P \cup Q}[j] \in S_k^q\}$, $j^+ = \min\{j > x \mid \text{SSA}_{P \cup Q}[j] \in S_k^q\}$, if they exist. Then, the

larger of $\text{lcp}(i, \text{SSA}_{P \cup Q}[j^-])$ and $\text{lcp}(i, \text{SSA}_{P \cup Q}[j^+])$ gives $\pi(i, S_k^q) = \arg \max_{i' \in S_k^q} \{\text{lcp}(i, i')\}$. Since $i, \text{SSA}_{P \cup Q}[j^+], \text{SSA}_{P \cup Q}[j^-] \in P \cup Q$, these values can be computed in constant time, resulting in a total of $O(\frac{n}{\tau} \log \tau)$ time for all i and k . Since the union of S_k^q over all q is S_k , we have $\pi(i, S_k) = \pi(i, S_k^{\hat{q}})$ and $L(i, S_k) = L(i, S_k^{\hat{q}})$, where $\hat{q} = \arg \max_{0 \leq q' < \tau} \{\text{lcp}(i, \pi(i, S_k^{q'}))\}$, so we can obtain $\pi(i, S_k)$ and $L(i, S_k)$ for each sampled position i and S_k by repeating the above process for each q , taking $O(n \log \tau)$ time. Thus, the total time for preprocessing, dominated by Corollary 8, is $O(n\tau)$.

► **Theorem 10.** *For any string of length n and integer $1 \leq \tau \leq \frac{n}{\log n}$, a data structure of size $O(\frac{n}{\tau})$ can be constructed in $O(n\tau)$ time using $O(\frac{n}{\tau})$ space, such that for any $0 \leq i, j < n$, $\text{lcp}(i, j)$ can be answered in $O(\tau \log \tau)$ time.*

4.4 Fast construction of combined data structure

The construction of the combined data structure is done using the same algorithms as described in Sections 4.2 and 4.3, with only minor modifications. For Bille et al.'s data structure, we only need to consider in addition to sampled positions, the positions in $D = \{i' \mid (i' + d) \bmod t = p, 0 \leq d < d_t\}$ due to the modification introduced for the combination. This reduces the array sizes (and thus the computation time) needed for the computation of $\pi(i, I_{\ell(v)})$ and $\pi(i, I_{r(v)})$ (and $L(i, I_{\ell(v)})$ and $L(i, I_{r(v)})$) to $O(\frac{n}{t} + \frac{n}{t} \cdot \frac{t^2}{n} \cdot \frac{1}{\tau}) = O(\frac{n}{t} + \frac{t}{\tau}) = O(\frac{n}{\tau \log \frac{n}{\tau}} + \log \frac{n}{\tau})$ for a total of $O(\frac{n}{\tau} + \log^2 \frac{n}{\tau}) = O(\frac{n}{\tau})$ for all levels, and for all q , we get $O(n)$ time. Thus, the total time for preprocessing is now dominated by Corollary 8, and is $O(n\tau)$.

► **Theorem 11.** *For any string of length n and integer $1 \leq \tau \leq n$, a data structure of size $O(\frac{n}{\tau})$ can be constructed in $O(n\tau)$ time using $O(\frac{n}{\tau})$ space, such that for any $0 \leq i, j < n$, $\text{lcp}(i, j)$ can be answered in $O(\tau \log \frac{n}{\tau})$ time.*

As noted previously, since $\tau \leq \frac{n}{\tau}$ when $\tau \leq \sqrt{n}$, and $\tau \geq \frac{n}{\tau}$ when $\tau \geq \sqrt{n}$, we get the following by simply choosing the data structure of Theorems 10 and 11, depending on the value of τ .

► **Corollary 12.** *For any string of length n and integer $1 \leq \tau \leq n$, a data structure of size $O(\frac{n}{\tau})$ can be constructed in $O(n\tau)$ time using $O(\frac{n}{\tau})$ space, such that for any $0 \leq i, j < n$, $\text{lcp}(i, j)$ can be answered in $O(\tau \min\{\log \tau, \log \frac{n}{\tau}\})$ time.*

5 Applications

Using the proposed data structure, the lexicographic order between two arbitrary suffixes can be computed in $O(\tau \min\{\log \tau, \log \frac{n}{\tau}\})$ time using $O(\frac{n}{\tau})$ space. Thus, using any $O(n \log n)$ comparison based sorting algorithm, we can compute the suffix array of a string of length n in $O(\min\{\log \tau, \log \frac{n}{\tau}\} n \tau \log n)$ time using $O(\frac{n}{\tau})$ working space, excluding the input and output. The best known deterministic space/time trade-off is $O(n\tau^2)$ time (for $1 \leq \tau \leq \sqrt[4]{n}$) using the same space [7], and our algorithm is better when $\tau = \Omega(\log^{1+\epsilon} n)$ for any $\epsilon > 0$.

Acknowledgements. The authors thank the anonymous reviewers for careful reading of the paper and for helpful comments.

References

- 1 Michael A. Bender and Martin Farach-Colton. The LCA problem revisited. In *Proc. Latin'00*, pages 88–94, 2000.
- 2 Philip Bille, Inge Li Gørtz, Mathias Bæk Tejs Knudsen, Moshe Lewenstein, and Hjalte Wedel Vildhøj. Longest common extensions in sublinear space. In *Proc. CPM 2015*, pages 65–76, 2015.
- 3 Philip Bille, Inge Li Gørtz, Benjamin Sach, and Hjalte Wedel Vildhøj. Time-space trade-offs for longest common extensions. *J. Discrete Algorithms*, 25:42–50, 2014.
- 4 Dan Gusfield. *Algorithms on Strings, Trees, and Sequences*. Cambridge University Press, 1997.
- 5 Juha Kärkkäinen. Fast BWT in small space by blockwise suffix sorting. *Theor. Comput. Sci.*, 387(3):249–257, 2007.
- 6 Juha Kärkkäinen, Dominik Kempa, and Simon J. Puglisi. Linear time Lempel-Ziv factorization: Simple, fast, small. In *Proc. CPM'13*, volume 7922 of *Lecture Notes in Computer Science*, pages 189–200. Springer, 2013.
- 7 Juha Kärkkäinen, Peter Sanders, and Stefan Burkhardt. Linear work suffix array construction. *J. ACM*, 53(6):918–936, 2006.
- 8 Juha Kärkkäinen and Esko Ukkonen. Sparse suffix trees. In *Proc. COCOON'96*, pages 219–230, 1996.
- 9 Toru Kasai, Gunho Lee, Hiroki Arimura, Setsuo Arikawa, and Kunsoo Park. Linear-time longest-common-prefix computation in suffix arrays and its applications. In *Proc. CPM'01*, pages 181–192, 2001.
- 10 Dong Kyue Kim, Jeong Seop Sim, Heejin Park, and Kunsoo Park. Linear-time construction of suffix arrays. In *Proc. CPM'03*, pages 186–199, 2003.
- 11 Pang Ko and Srinivas Aluru. Space efficient linear time construction of suffix arrays. In *Proc. CPM'03*, pages 200–210, 2003.
- 12 Udi Manber and Gene Myers. Suffix arrays: A new method for on-line string searches. *SIAM J. Comput.*, 22(5):935–948, 1993.
- 13 Simon J. Puglisi and Andrew Turpin. Space-time tradeoffs for longest-common-prefix array computation. In *Proc. ISAAC'08*, volume 5369 of *Lecture Notes in Computer Science*, pages 124–135. Springer, 2008.