

Weighted Ancestors in Suffix Trees Revisited

Djamal Belazzougui ✉

Centre de Recherche sur l'Information Scientifique et Technique, Algiers, Algeria

Dmitry Kosolobov ✉ 

Institute of Natural Sciences and Mathematics, Ural Federal University, Ekaterinburg, Russia

Simon J. Puglisi ✉ 

Department of Computer Science, University of Helsinki, Helsinki, Finland

Rajeev Raman ✉ 

Department of Informatics, University of Leicester, Leicester, United Kingdom

Abstract

The weighted ancestor problem is a well-known generalization of the predecessor problem to trees. It is known to require $\Omega(\log \log n)$ time for queries provided $\mathcal{O}(n \text{ polylog } n)$ space is available and weights are from $[0..n]$, where n is the number of tree nodes. However, when applied to suffix trees, the problem, surprisingly, admits an $\mathcal{O}(n)$ -space solution with constant query time, as was shown by Gawrychowski, Lewenstein, and Nicholson (Proc. ESA 2014). This variant of the problem can be reformulated as follows: given the suffix tree of a string s , we need a data structure that can locate in the tree any substring $s[p..q]$ of s in $\mathcal{O}(1)$ time (as if one descended from the root reading $s[p..q]$ along the way). Unfortunately, the data structure of Gawrychowski et al. has no efficient construction algorithm, limiting its wider usage as an algorithmic tool. In this paper we resolve this issue, describing a data structure for weighted ancestors in suffix trees with constant query time and a linear construction algorithm. Our solution is based on a novel approach using so-called irreducible LCP values.

2012 ACM Subject Classification Theory of computation \rightarrow Pattern matching

Keywords and phrases suffix tree, weighted ancestors, irreducible LCP, deterministic substring hashing

Digital Object Identifier 10.4230/LIPIcs.CPM.2021.8

Funding *Dmitry Kosolobov*: Supported by the grant 20-71-00050 of the Russian Foundation of Basic Research (for the final version of the data structure and for its construction algorithm).

1 Introduction

The suffix tree is one of the central data structures in stringology. Its primary application is to determine whether an arbitrary string t occurs as a substring of another string, s , which can be done in time proportional to the length of t by traversing the suffix tree of s downward from the root and reading off the symbols of t along the way. Many algorithms using suffix trees perform this procedure for substrings $t = s[p..q]$ of s itself. In this important special case, the traversal can be executed much faster than $\mathcal{O}(q - p + 1)$ time provided the tree has been preprocessed to build some additional data structures [1, 11, 13]; particularly surprising is that the traversal can be performed in constant time using only linear space, as shown by Gawrychowski et al. [13]. In this paper, we describe the first linear construction algorithm for such a data structure. The lack of an efficient construction algorithm for the result of [13] has been, apparently, the main obstacle hindering its wider adoption. Our solution is completely different from that of [13] and is based on a combinatorial result of Kärkkäinen et al. [17] (see also [16]) that estimates the sum of irreducible LCP values (precise definitions follow).



© Djamal Belazzougui, Dmitry Kosolobov, Simon J. Puglisi, and Rajeev Raman; licensed under Creative Commons License CC-BY 4.0

32nd Annual Symposium on Combinatorial Pattern Matching (CPM 2021).

Editors: Paweł Gawrychowski and Tatiana Starikovskaya; Article No. 8; pp. 8:1–8:15

Leibniz International Proceedings in Informatics



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

As one might expect, the described data structure has a multitude of applications: [1, 7, 8, 11, 19, 20, 21], to name a few. We, however, do not dive further into a discussion of these applications and refer the reader to the overview in [13, Sect. 1] and references therein for more details.

In order to “locate” a substring $t = s[p..q]$ in the suffix tree of s , it suffices to answer the following query: given a node of the tree that corresponds to the suffix of s starting at position p (usually, it is a leaf), we should find the farthest (closest to the root) ancestor v of this node such that the string written on the root– v path has length at least $q - p + 1$. This problem is a particular case of the general *weighted ancestor problem* [1, 11, 22]: given a tree whose nodes are associated with integer weights such that the weights decrease if one ascends from any node to the root (the weight of a node in the suffix tree is the length of the string written on the root–node path), the tree should be preprocessed in order to answer *weighted ancestor queries* that, for a given node v and a number w , return the farthest ancestor of v whose weight is at least w . The problem can be viewed as a generalization to trees of the predecessor search problem, in which we preprocess a set of integers to support predecessor queries: for any given number w , return the largest integer from the set that precedes w .

Clearly, any linear-space solution for weighted ancestors can be used as a solution for predecessor search. As was shown in [22] and [13], a certain converse reduction is also possible: the weighted ancestor queries for a tree with n nodes and integer weights from a range $[0..U]$ can be answered in $\mathcal{O}(\text{pred}(n, U))$ time using linear space, where $\text{pred}(n, U)$ is the time required to answer predecessor search queries for any set of $k \leq n$ integers from the range $[0..U]$ using $\mathcal{O}(k)$ space. Therefore, when $U = n$ (as in the case of suffix tree), both problems can be solved in linear space with $\mathcal{O}(\log \log n)$ -time queries using the standard van Emde Boas or y-fast trie data structures [30, 31].

Due to the lower bound of Pătraşcu and Thorup [25], the time $\mathcal{O}(\log \log n)$ is optimal for $U = n$ when the available space is linear, and moreover, any solution of the weighted ancestor problem that uses $\mathcal{O}(n \text{ polylog } n)$ space must spend $\Omega(\log \log n)$ time on queries (see [13, Appendix A]). In view of this lower bound, it is all the more unexpected that the special case of suffix trees admits an $\mathcal{O}(n)$ -space solution with constant query time. In order to circumvent the lower bound, Gawrychowski et al. [13] solve predecessor search problems on some paths of the suffix tree using $\mathcal{O}(n)$ bits of space (or slightly less), which admits a constant time solution by a so-called rank data structure; because of the internal repetitive structure of the suffix tree, the solution for one path can be reused in many different paths in such a way that, in total, the utilized space is linear. Our approach essentially relies on the same intuition but we perform path predecessor queries on different trees closely related to the suffix tree and the advantages of repetitive structures come implicitly; in particular, we do not explicitly treat periodic regions of the string separately so that, in this regard, our solution is more “uniform”, in a sense, than that of [13].

This paper is organized as follows. In Section 2 the problem is reduced to certain path-counting queries using (unweighted) level ancestor queries. Section 3 describes a simple solution for the queries locating substrings $s[p..q]$ whose position p corresponds to an irreducible LCP value. In Section 4 we reduce general queries to the queries at irreducible positions and a certain geometric orthogonal predecessor problem. Section 5 describes a solution for this special geometric problem. Conclusions and reflections are then offered in Section 6.

2 Basic Data Structures

Let us fix a string s of length n , denoting its letters by $s[0], s[1], \dots, s[n-1]$. We write $s[p..q]$ for the *substring* $s[p]s[p+1] \cdots s[q]$, assuming it is empty if $p > q$; $s[p..q]$ is called a *suffix* (resp., *prefix*) of s if $q = n - 1$ (resp., $p = 0$). For any string t , let $|t|$ denote its length. We say that t *occurs* at position p in s if $s[p..p+|t|-1] = t$. Denote $[p..q] = \{k \in \mathbb{Z}: p \leq k \leq q\}$.

A *suffix tree* of s is a compacted trie containing all suffixes of s . The labels on the edges of the tree are stored as pointers to corresponding substrings of s . For each tree node v , denote by $\text{str}(v)$ the string written on the root– v path. The number $|\text{str}(v)|$ is called the *string depth* of v . The *locus* of a substring $s[p..q]$ is the (unique) node v of minimal string depth such that $s[p..q]$ is a prefix of $\text{str}(v)$.

The string s and its suffix tree T are the input to our algorithm. To simplify the exposition, we assume that $s[n-1]$ is equal to a special letter $\$$ that is smaller than all other letters in s , so that there is a one-to-one correspondence between the suffixes of s and the leaves of T . Our computational model is the word-RAM and space is measured in $\Theta(\log n)$ -bit machine words. Our goal is to construct in $\mathcal{O}(n)$ time a data structure that can find in the tree the locus of any substring $s[p..q]$ of s in $\mathcal{O}(1)$ time.

A *level ancestor query* in a tree asks, for a given node v and an integer $d \geq 1$, the d th node on the v –root path (provided the path has at least d nodes). It is known that any tree can be preprocessed in linear time to answer such queries in constant time [5, 6]. With such a structure, the locus of a substring $s[p..q]$ can be found by first locating the leaf v of T corresponding to $s[p..n-1]$ (i.e., $\text{str}(v) = s[p..n-1]$), which can be located using a precomputed array of length n and, then, counting the number d of nodes u on the v –root path such that $|\text{str}(u)| > q - p$; then, evidently, the locus of $s[p..q]$ is given by the level ancestor query on v and d .

We have to complicate this scheme slightly since the machinery that we develop in the sequel allows us to count only those nodes u on the v –root path that have a branch to the “left” of the path (or, symmetrically, to the “right”). More formally, given a node v , a node u on the v –root path is called *left-branching* (resp., *right-branching*) if there is a suffix $s[p..n-1]$ that is lexicographically smaller (resp., greater) than the string $\text{str}(v)$ and its longest common prefix with $\text{str}(v)$ is $\text{str}(u)$. For instance, the path from the leaf corresponding to the string *sippi* $\$$ in Figure 1 has four nodes but only one of them (namely, the root) is left-branching.


► **Lemma 1.** *For any suffix tree, one can build in linear time a data structure that, for any node v and integer $d \geq 1$, can return in $\mathcal{O}(1)$ time the d th left-branching (or right-branching) node on the v –root path.*

Proof. Traversing the suffix tree, we construct another tree on the same set of nodes in which the parent of each non-root node is either its nearest left-branching ancestor in the suffix tree (if any) or the root. The queries for left-branching nodes can be answered by the level ancestor structure [5, 6] built on this new tree. The right-branching case is symmetric. ◀

Thus, to find the locus for $s[p..q]$, we have to count on a leaf–root path the number of left- and right-branching nodes whose string depths are greater than the threshold $q - p$; we then use these two numbers in the data structure of Lemma 1 in order to find two candidate nodes and the node with smaller string depth is the locus. In the remaining text, we focus only on this counting problem and only on left-branching nodes as the right-branching case is symmetric. First, however, we define a number of useful standard structures.

The *suffix array* of s is an array $\text{sa}[0..n-1]$ containing integers from 0 to $n - 1$ such that $s[\text{sa}[0]..n-1] < s[\text{sa}[1]..n-1] < \cdots < s[\text{sa}[n-1]..n-1]$ lexicographically [23]. The *inverse suffix array*, denoted $\text{isa}[0..n-1]$, is defined as $\text{sa}[\text{isa}[p]] = p$, for all $p \in [0..n-1]$. For any

positions p and q , denote by $\text{lcp}(p, q)$ the length of the longest common prefix of $s[p..n-1]$ and $s[q..n-1]$. The *longest common prefix (LCP) array* is an array $\text{lcp}[0..n-1]$ such that $\text{lcp}[0] = 0$ and $\text{lcp}[i] = \text{lcp}(\text{sa}[i-1], \text{sa}[i])$, for $i \in [1..n-1]$. The *permuted longest common prefix (PLCP) array* is an array $\text{plcp}[0..n-1]$ such that $\text{plcp}[p] = \text{lcp}[\text{isa}[p]]$, for $p \in [0..n-1]$. The *Burrows–Wheeler transform* [9] is a string $\text{bwt}[0..n-1]$ such that $\text{bwt}[i] = s[\text{sa}[i]-1]$ if $\text{sa}[i] \neq 0$, and $\text{bwt}[i] = s[n-1]$ otherwise. All the arrays sa , isa , lcp , plcp and the string bwt can be built from T in $\mathcal{O}(n)$ time. Some of these structures are depicted in Figure 1.



■ **Figure 1** The suffix tree T , lcp , sa , and bwt of the string $s = \text{mississippi}\$$. All irreducible LCP values are in bold and their sum is 7; all positions of s except 2, 3, 4 are irreducible.

3 Queries at Irreducible Positions

An LCP value $\text{lcp}[i]$ is called *irreducible* if either $i = 0$ or $\text{bwt}[i-1] \neq \text{bwt}[i]$. In other words, the irreducible LCP values are those values of $\text{lcp}[0..n-1]$ that correspond to the first positions in the runs of bwt (e.g., all values $\text{lcp}[i]$ except for $i = 3, 9, 11$ in Figure 1). The central combinatorial fact that is at the core of our construction is the following surprising lemma proved by Kärkkäinen et al. [17].

► **Lemma 2** (see [17]). *The sum of all irreducible LCP values is at most $2n \log n$.*

We say that a position p of the string s is *irreducible* if $\text{lcp}[\text{isa}[p]] (= \text{plcp}[p])$ is an irreducible LCP value.

Consider a query that asks to find the locus of a substring $s[p..q]$ such that p is an irreducible position. Denote $\ell = q - p + 1$. We first locate in $\mathcal{O}(1)$ time the leaf v corresponding to $s[p..n-1]$ using a precomputed array. As was discussed above, the problem is reduced to the counting of left-branching nodes u on the v -root path such that $|\text{str}(u)| \geq \ell$. We will construct for p a bit array $b_p[0..\ell_p-1]$ of length $\ell_p = \text{plcp}[p]$ such that, for any d , $b_p[d] = 1$ iff there is a left-branching node with string depth d on the v -root path. Since the suffix $s[p..n-1]$ and its lexicographical predecessor among all suffixes (if any, i.e., if $\text{isa}[p] \neq 0$) have the longest common prefix of length $\text{plcp}[p]$, the lowest left-branching node on the v -root path has string depth $\text{plcp}[p]$. Therefore, the number of left-branching nodes u on the v -root path such that $|\text{str}(u)| \geq \ell$ is equal to the number of ones in the subarray $b_p[\ell..\ell_p-1]$ plus 1. The bit counting query is answered using the following *rank data structure*.

► **Lemma 3** (see [10, 15]). *For any bit array $b[0..m-1]$ packed into $\mathcal{O}(m/w)$ w -bit machine words, one can construct in $\mathcal{O}(m/w)$ time a rank data structure that can count the number of ones in any range $b[p..q]$ in $\mathcal{O}(1)$ time, provided a table computable in $o(2^w)$ time and independent of the array has been precomputed.*

By Lemma 2, the total length of all arrays b_p for all irreducible positions p of s is $\mathcal{O}(n \log n)$. Therefore, one can construct the rank data structures of Lemma 3 for them in $\mathcal{O}(n)$ overall time provided $w = \lfloor \log n \rfloor$. It remains to build the arrays b_p themselves.

To this end, we traverse the suffix tree T in depth first (i.e., lexicographical) order, maintaining along the way a bit array b of length n such that, when we are in a node v , the subarray $b[0..|\text{str}(v)|-1]$ marks the string depths of all left-branching ancestors of v , i.e., for $0 \leq d < |\text{str}(v)|$, we have $b[d] = 1$ iff there is a left-branching node with string depth d on the v -root path. The array is stored in a packed form in $\lfloor \log n \rfloor$ -bit machine words. For each node v , we consider its outgoing edges. Each pair of adjacent edges (in the lexicographical order of their labels) corresponds to a unique value $\text{lcp}[i]$ such that $s[\text{sa}[i]..n-1]$ is the suffix corresponding to the leftmost (lexicographically smallest) leaf of the subtree connected to the larger one of the two edge labels. We check whether the value $\text{lcp}[i]$ is irreducible comparing $\text{bwt}[i-1]$ and $\text{bwt}[i]$ and, if so, store the subarray $b[0..|\text{str}(v)|-1]$ into $b_{\text{sa}[i]}$ in $\mathcal{O}(1 + |\text{str}(v)|/\log n)$ time. As we touch in this way every value $\text{lcp}[i]$ only once, the overall time is $\mathcal{O}(n)$ by the same argument using Lemma 2.

4 Reduction to Irreducible Positions

Consider a query for the locus of a substring $s[p..q]$ such that the position p is not irreducible. Let v be the leaf corresponding to $s[p..n-1]$. As was discussed in Section 2, to answer the query, we have to count the number of left-branching nodes with string depths at least $q - p + 1$ on the v -root path. As a first approach, we do the following precalculations for this.

We build in $\mathcal{O}(n)$ time on the suffix tree T a data structure that allows us to find the lowest common ancestor for any pair of nodes in $\mathcal{O}(1)$ time [4, 14]. In one tree traversal, we precompute in each node u the number of left-branching nodes on the u -root path. We create two arrays $R_{\leq}[0..n-1]$ and $R_{\geq}[0..n-1]$ such that, for any i , $R_{\leq}[i] = \max\{j \leq i : j = 0 \text{ or } \text{bwt}[j-1] \neq \text{bwt}[j]\}$ and $R_{\geq}[i] = \min\{j \geq i : j = 0 \text{ or } \text{bwt}[j-1] \neq \text{bwt}[j]\}$ ($R_{\geq}[i]$ is undefined if there is no such j), i.e., $R_{\leq}[i]$ and $R_{\geq}[i]$ are indexes of irreducible LCP values, respectively, preceding and succeeding $\text{lcp}[i]$. Thus, for any suffix $s[p..n-1]$, the suffixes $s[\text{sa}[R_{\leq}[\text{isa}[p]]]..n-1]$ and $s[\text{sa}[R_{\geq}[\text{isa}[p]]]..n-1]$ are, respectively, the closest lexicographical predecessor and successor of $s[p..n-1]$ with irreducible starting position. The *closest irreducible lexicographical neighbour* of $s[p..n-1]$ is that suffix $s[r..n-1]$ among these two that has longer common prefix with $s[p..n-1]$, i.e., $r = \text{sa}[R_{\leq}[\text{isa}[p]]]$ if either $\text{lcp}(\text{sa}[R_{\leq}[\text{isa}[p]]], p) \geq \text{lcp}(\text{sa}[R_{\geq}[\text{isa}[p]]], p)$ or $R_{\geq}[\text{isa}[p]]$ is undefined, and $r = \text{sa}[R_{\geq}[\text{isa}[p]]]$ otherwise. Since $\text{lcp}(t, p)$ can be calculated in $\mathcal{O}(1)$ time for any positions t and p by one lowest common ancestor query on their corresponding leaves, the closest irreducible lexicographical neighbour can be computed in $\mathcal{O}(1)$ time.

Now, consider a query for the locus of $s[p..q]$ with non-irreducible p . We first find in $\mathcal{O}(1)$ time the closest irreducible lexicographical neighbour $s[r..n-1]$ for $s[p..n-1]$. Let v' and v be the leaves corresponding to $s[r..n-1]$ and $s[p..n-1]$, respectively. We compute in $\mathcal{O}(1)$ time the lowest common ancestor u of v' and v . Thus, $|\text{str}(u)| = \text{lcp}(r, p)$. Using the number of left-branching nodes on the v -root and u -root paths that were precomputed in v and u , we calculate in $\mathcal{O}(1)$ time the number k of left-branching nodes on the v -root path that lie between the nodes v and u (inclusively).

Denote $\ell = q - p + 1$. If $\ell \leq |\text{str}(u)|$, then we can count in the v -root path the number of left-branching nodes w such that $|\text{str}(w)| \geq \ell$ as follows. The number of nodes w such that $|\text{str}(w)| \geq |\text{str}(u)|$ is equal to k . The number of nodes w such that $\ell \leq |\text{str}(w)| < |\text{str}(u)|$ can be found by counting the number of ones in the subarray $b_r[\ell..|\text{str}(u)|-1]$ of the bit

array b_r associated with the irreducible position r (assuming that all values $b_r[t]$ with $t \geq \ell_r$ are zeros in case the length ℓ_r of b_r is less than $|\text{str}(u)|$), which can be performed in $\mathcal{O}(1)$ time by Lemma 3. Thus, the problem is solved for the case $\ell \leq |\text{str}(u)|$.

To address the case $\ell > |\text{str}(u)|$, we have to develop more sophisticated techniques that allow us to reduce the counting of left-branching nodes on the v -root path to counting on a different leaf-root path with a different threshold ℓ' that meets the condition $\ell' \leq |\text{str}(u')|$, for an analogously appropriately defined node u' . The remainder of the text describes the reduction.


Similar to irreducible positions, let us define, for each non-irreducible position p , a number $\ell_p = \text{plcp}[p]$ and an array $b_p[0..\ell_p-1]$ such that, for any d , $b_p[d] = 1$ iff there is a left-branching node of string depth d on the path from the leaf corresponding to $s[p..n-1]$ to the root. We do not actually store the arrays b_p and use them only in the analysis.

► **Lemma 4.** *For any non-irreducible position p , we have $\ell_{p-1} = \ell_p + 1$ and, if $b_{p-1}[d+1] = 1$ for some $d \geq 0$, then $b_p[d] = 1$.*

Proof. Let $s[t..n-1]$ be the suffix lexicographically preceding $s[p..n-1]$, i.e., $t = \text{sa}[\text{isa}[p]-1]$ and $\ell_p = \text{lcp}(t, p)$. Since p is not irreducible, we have $\text{bwt}[\text{isa}[t]] = \text{bwt}[\text{isa}[p]]$, i.e., $s[t-1] = s[p-1]$. Hence, $s[t-1..n-1]$ is the suffix lexicographically preceding $s[p-1..n-1]$: $\text{isa}[t-1] = \text{isa}[p-1] - 1$. Therefore, $\ell_{p-1} = \text{lcp}(t-1, p-1)$, which is equal to $\text{lcp}(t, p) + 1 = \ell_p + 1$.

If $b_{p-1}[d+1] = 1$ for $d \geq 0$, then there is a suffix $s[r..n-1]$ that is smaller than $s[p-1..n-1]$ and $\text{lcp}(r, p-1) = d+1$. Then, $\text{lcp}(r+1, p) = d$, which implies that $b_p[d] = 1$. ◀

Consider two consecutive irreducible positions p' and p'' in s (i.e., all positions r between p' and p'' are not irreducible). Lemma 4 states that the arrays $b_{p'}, b_{p'+1}, \dots, b_{p''-1}$ form a trapezoidal structure as depicted in Figure 2 in which each 1 value is inherited by all arrays above it while it fits in their range.




■ **Figure 2** Here, p' and p'' are consecutive irreducible positions. Each line of the trapezoid depicts an array $b_p[0..\ell_p-1]$, for $p \in [p'..p''-1]$, in which gray and white positions signify, respectively, ones and zeros.

The query for the locus of $s[p..q]$ was essentially reduced to the counting of ones in the subarray $b_p[\ell..\ell_p-1]$, where $\ell = q - p + 1$. The problem now is that the array b_p is not stored explicitly since the position p is not irreducible. Denote the length of $b_p[\ell..\ell_p-1]$ by $m = \ell_p - \ell$, so that $b_p[\ell..\ell_p-1] = b_p[\ell_p-m..\ell_p-1]$, which is a more convenient notation for what follows. Let p' be the closest irreducible position preceding p , i.e., $p' = \max\{r \leq p : r \text{ is irreducible}\}$. If we are lucky and neither of the arrays $b_{p'+1}, b_{p'+2}, \dots, b_p$ introduces new 1 values in the subarray $b_{p'}[\ell_{p'}-m..\ell_{p'}-1]$ (in other words if $b_{p'}[\ell_{p'}-m..\ell_{p'}-1] = b_p[\ell_p-m..\ell_p-1]$), then we can simply count the number of ones in $b_{p'}[\ell_{p'}-m..\ell_{p'}-1]$ in $\mathcal{O}(1)$ time using Lemma 3 and, thus, solve the problem. Unfortunately, new 1 values indeed could be introduced. But, as it turns out, such new values are, in a sense, “covered” by other arrays b_r at some irreducible positions r as the following lemma suggests.

► **Lemma 5.** *Let p be a non-irreducible position and $s[r..n-1]$ be the closest irreducible lexicographical neighbour of $s[p..n-1]$. Denote $c_p = \text{lcp}(r, p)$. If, for some d , we have $b_p[d] = 1$ but $b_{p-1}[d+1] = 0$, then $d \leq c_p$.*

Proof. The condition $b_p[d] = 1$ implies that there is a suffix $s[t..n-1]$ that is smaller than $s[p..n-1]$ and $\text{lcp}(t, p) = d$. If $\text{bwt}[\text{isa}[t]] = \text{bwt}[\text{isa}[p]]$, then $s[t-1] = s[p-1]$ and, hence, the suffix $s[t-1..n-1]$ is smaller than $s[p-1..n-1]$ and $\text{lcp}(t-1, p-1) = d+1$. This gives $b_{p-1}[d+1] = 1$, which contradicts the equality $b_{p-1}[d+1] = 0$. Hence, $\text{bwt}[\text{isa}[t]] \neq \text{bwt}[\text{isa}[p]]$. Thus, at least one of the values $\text{lcp}[\text{isa}[t]+1], \text{lcp}[\text{isa}[t]+2], \dots, \text{lcp}[\text{isa}[p]-1]$ must be irreducible (note that $\text{isa}[t]+1 < \text{isa}[p]$ since otherwise p would be irreducible). Let r' be an irreducible position such that $\text{isa}[t] < \text{isa}[r'] < \text{isa}[p]$. We obtain $\text{lcp}(r, p) \geq d$ since $\text{lcp}(r', p) \geq \text{lcp}(t, p) = d$ and, for any irreducible r'' (in particular, for r'), we have $\text{lcp}(r, p) \geq \text{lcp}(r'', p)$. ◀

Observe that c_p in Lemma 5 is equal to $|\text{str}(u)|$, where u is the lowest common ancestor of the leaves v' and v corresponding to $s[r..n-1]$ and $s[p..n-1]$. Thus, the numbers c_p can be precomputed for all non-irreducible positions p in $\mathcal{O}(n)$ time using lowest common ancestor queries and the arrays R_{\leq} and R_{\geq} described in the beginning of the present section. For irreducible positions p' , we put $c_{p'} = \ell_{p'}$ by definition. To illustrate Lemma 5 and its consequences, we depict in Figure 3 the same trapezoidal structure as in Figure 2 coloring in each b_p a prefix of length $c_p + 1$ (and also depicting the range $b_p[\ell_p - m.. \ell_p - 1]$).



■ **Figure 3** The yellow stripe in each array b_p is a prefix of length $\min\{c_p + 1, \ell_p\}$. Note that, according to Lemma 5, each gray position whose corresponding position below is white is covered in yellow. The blue stripe is $b_p[\ell_p - m.. \ell_p - 1]$, where $p = p' + 4$. It coincides with two corresponding regions below it that are in light blue. Here we have $t = p - 2$.

The problematic condition $\ell > |\text{str}(u)|$ that we consider can be reformulated as $\ell_p - m > c_p$. It follows from Lemmas 4 and 5 that in this case $b_p[\ell_p - m.. \ell_p - 1] = b_{p-1}[\ell_{p-1} - m.. \ell_{p-1} - 1]$. Let t be the first position preceding p such that $\ell_t - m \leq c_t$. Note that $t \geq p'$ since $c_{p'} = \ell_{p'}$ by definition and, thus, $\ell_{p'} - m \leq c_{p'}$. Then, applying Lemma 5 consecutively to all positions $p, p-1, \dots, t+1$, we conclude that $b_t[\ell_t - m.. \ell_t - 1] = b_p[\ell_p - m.. \ell_p - 1]$.

Informally, in terms of Figure 3, the searching of t corresponds to the moving of the range $b_p[\ell_p - m.. \ell_p - 1]$ down until we encounter an “obstacle”, a colored part of b_t of length $c_t + 1$. The specific algorithm finding t is discussed in Section 5; let us assume for the time being that t is already known. Then, Lemma 5 implies that all the ranges $b_r[\ell_r - m.. \ell_r - 1]$ with $r \in [t..p]$ coincide and, thus, the whole problem was reduced to the counting of ones in the subarray $b_t[\ell_t - m.. \ell_t - 1]$. But since $\ell_t - m \leq c_t$, the problem can be solved by the method described at the beginning of the section: we find an irreducible position r such that $\text{lcp}(r, t) = c_t$, then locate the leaves v' and v'' corresponding to $s[r..n-1]$ and $s[t..n-1]$, respectively, find the lowest common ancestor u' of v' and v'' , and separately count the number of left-branching nodes w on the v'' -root path such that $|\text{str}(w)| \geq c_t$ and such that $\ell_t - m \leq |\text{str}(w)| < c_t$.


Let us briefly recap the reductions described above: the searching for the locus of $s[p..q]$ was essentially reduced to the counting of ones in the subarray $b_p[\ell_p - m.. \ell_p - 1]$, where $m = \ell_p - (q - p + 1)$, for which we first compute the position $t = \max\{t \leq p: \ell_t - m \leq c_t\}$ (this

is discussed in Section 5) and, then, count the number of ones in the subarray $b_t[\ell_t - m .. \ell_t - 1]$ (the subarray coincides with $b_p[\ell_p - m .. \ell_p - 1]$ by Lemmas 4 and 5) using lowest common ancestor queries, the arrays R_{\leq} and R_{\geq} , and some precomputed numbers in the nodes.

5 Reduction to Special Weighted Ancestors

We are to compute $t = \max\{t \leq p: \ell_t - m \leq c_t\}$, for a non-irreducible position p . As is seen in Figure 3, this is a kind of geometric “orthogonal range predecessor” problem.

For each irreducible position p' in s , we build a tree $I_{p'}$ with $p'' - p'$ nodes, where p'' is the next irreducible position after p' (so that all r with $p' < r < p''$ are not irreducible): Each node of $I_{p'}$ is associated with a unique position from $[p' .. p'' - 1]$ and the root is associated with p' . A node associated with position $r \neq p'$ has weight $w_r = r - p' + c_r$; the root has weight $w_{p'} = +\infty$. The parent of a non-root node associated with r is the node associated with the position $r' = \max\{r' < r: w_{r'} > w_r\}$. Thus, the weights strictly increase as one ascends to the root. The tree $I_{p'}$ is easier to explain in terms of the trapezoidal structure drawn in Figure 4: the parent of a node associated with r is its closest predecessor r' whose colored range $b_{r'}[r' .. r' + c_{r'} - 1]$ goes at least one position farther to the right than $b_r[r .. r + c_r - 1]$. By Lemma 4, $r - p' + \ell_r = \ell_{p'}$ and, hence, each weight w_r such that $c_r \leq \ell_r$ is upperbounded by $\ell_{p'}$. However, it may happen that $c_r > \ell_r$ and, thus, $w_r > \ell_{p'}$. In this case, the colored range $b_r[r .. r + c_r - 1]$ stretches beyond the length ℓ_r of b_r (for simplicity, Figure 4 does not have such cases). Such large weights are a source of technical complications in our scheme.



■ **Figure 4** Each stripe of the trapezoid depicts an array b_r , for $r \in [p' .. p'' - 1]$, and its prefix of length c_r (possibly empty) is colored in yellow (not $\min\{c_r + 1, \ell_r\}$ like in Figure 3). Each yellow prefix corresponds to a node of $I_{p'}$ and it is connected to the yellow prefix corresponding to its parent in $I_{p'}$. To avoid overloading the picture, gray positions signifying 1s in b_r are not drawn.

The tree can be constructed in linear time inserting consecutively the nodes associated with $p', p' + 1, \dots$ and maintaining a stack that contains all nodes of the path from the last inserted node to the root: to insert a new node with weight w , we pop from the stack all nodes until a node heavier than w is encountered to which the new node is attached.

In terms of the tree $I_{p'}$, the node associated with the sought position t is the nearest ancestor of the node associated with p such that $w_t \geq \ell + p - p'$, where $\ell = q - p + 1$: the condition $w_t \geq \ell + p - p'$ is equivalent to $\ell_t - m \leq c_t$ since $w_t = t - p' + c_t$, $m = \ell_p - \ell$, and $\ell_t - \ell_p = p - t$. Since the threshold $\ell + p - p'$ in the condition does not exceed $\ell_{p'}$, we will be able to ignore differences between weights larger than $\ell_{p'}$ in our algorithm by “cutting” them in a way. Still, even with this restriction, at first glance this looks like quite a general weighted ancestor problem that admits no constant time solution in the space available. However, we will be able to use a structure common to such trees in order to speed up the computation. The idea is to decompose the tree into heavy paths (see below), as is usually done for weighted ancestor queries, and perform fast queries on the paths via the use of more memory; the trick is that, with some care, this space can be shared among “similar” trees and can be “traded” for irreducible LCP values relying again on Lemma 2.

Consider the tree $I_{p'}$. An edge (v, u) connecting a child v to its parent u is called *heavy* if $\text{size}(v) > \text{size}(u)/2$, and *light* otherwise, where $\text{size}(w)$ denotes the number of nodes in the subtree rooted at w . Thus, at most one child can be connected to u by a heavy edge. One can easily show that any leaf–root path contains at most $\log n$ light edges. All nodes of $I_{p'}$ are decomposed into disjoint *heavy paths*, maximal paths with only heavy edges in them. Note that in this version of the heavy-light decomposition [28] the number of heavy edges incident to a given non-leaf node can be zero (it then forms a singleton path). The decomposition can be constructed in linear time in one tree traversal.

The *predecessor search problem*, for a set of increasing numbers w_1, w_2, \dots, w_k and a given threshold $w \leq w_k$, is to find $\min\{w_i : w \leq w_i\}$. Although the problem, in fact, searches for a successor, we call it “predecessor search” as it is essentially an equivalent problem.

► **Lemma 6** ([13, Lem. 11]). *Consider a tree on m nodes with weights from $[0..n]$ such that some of the nodes are marked, any leaf–root path contains at most $\mathcal{O}(\log n)$ marked nodes, and the weights on any leaf–root path increase. One can preprocess the tree in $\mathcal{O}(m)$ time so that predecessor search can be performed among the marked ancestors of any node in $\mathcal{O}(1)$ time, provided a table computable in $o(n)$ and independent of the tree is precalculated.¹*

For each heavy path, we mark the node closest to the root. Then, the data structure of Lemma 6 is built on each tree $I_{p'}$, which takes $\mathcal{O}(n)$ total time for all the trees $I_{p'}$. To answer a weighted ancestor query on $I_{p'}$, we find in $\mathcal{O}(1)$ time the heavy path containing the answer using this data structure and, then, consider the predecessor problem inside the path.

Consider a heavy path whose node weights are $w_{i_1}, w_{i_2}, \dots, w_{i_k}$ in increasing order. We have to answer a predecessor query on the path for a threshold w such that $w \leq \ell_{p'}$ (recall that only thresholds $\leq \ell_{p'}$ are of interest for us) and it is known that its result is in the path, i.e., $w \leq w_{i_k}$. Let $w_{i_m} = \max\{w_{i_j} : w_{i_j} \leq \ell_{p'}\}$. A trivial constant-time solution for such queries is to construct a bit array $a[0..c]$, where $c = w_{i_m} - w_{i_1}$, endowed with a rank data structure such that, for any d , we have $a[d] = 1$ iff $d = w_{i_j} - w_{i_1}$, for some $j \in [1..m]$. Then, the predecessor query with a threshold w such that $w \leq \min\{\ell_{p'}, w_{i_k}\}$ can be answered by counting the number h of 1s in the subarray $a[0..w - w_{i_1} - 1]$, thus giving the result $w_{i_{h+1}} = \min\{w_{i_j} : w \leq w_{i_j}\}$. The array a occupies $\mathcal{O}(1 + \ell_{p'} / \log n)$ space since $c \leq \ell_{p'}$.

It turns out that, with minor changes, the arrays a can be shared among many heavy paths in different trees. However, this approach *per se* leads to $\mathcal{O}(n \log n)$ time and space, as will be evident in the sequel. We therefore need a slightly more elaborate solution.

Instead of the array $a[0..c]$, we construct a bit array $\hat{a}[0.. \lfloor c / \lfloor \log n \rfloor \rfloor]$ endowed with a rank data structure such that, for any d , $\hat{a}[d] = 1$ iff the subarray $a[d \lfloor \log n \rfloor .. (d+1) \lfloor \log n \rfloor - 1]$ contains non-zero values (assuming that $a[i] = 0$, for $i > c$). The bit array \hat{a} packed into $\mathcal{O}(1 + c / \log^2 n)$ machine words of size $\lfloor \log n \rfloor$ bits can be built from the numbers $w_{i_1}, w_{i_2}, \dots, w_{i_k}$ in one pass in $\mathcal{O}(k + c / \log^2 n)$ time. For each 1 value, $\hat{a}[d] = 1$, we collect all weights w_{i_j} such that $d \lfloor \log n \rfloor \leq w_{i_j} - w_{i_1} < (d+1) \lfloor \log n \rfloor$ into a set S_d . The sets S_d , for all d such that $\hat{a}[d] = 1$, are disjoint and non-empty, and can be assembled in one pass through the weights in $\mathcal{O}(k)$ time. We also store pointers P_h , with $h = 1, 2, \dots$ ($h \leq m$), such that P_h refers to a non-empty set S_d such that $\hat{a}[d] = 1$ is the h th 1 value in \hat{a} (i.e., h is the number of 1s in the subarray $\hat{a}[0..d]$). Each set S_d is equipped with the following *fusion heap* data structure.

► **Lemma 7** (see [12, 26]). *For any set S of $\mathcal{O}(\log n)$ integers from $[0..n]$, one can build in $\mathcal{O}(|S|)$ time a fusion heap that answers predecessor queries in $\mathcal{O}(1)$ time.*

¹ Although Lemma 11 in [13] does not claim linear construction time, it easily follows from its proof.

With this machinery, a predecessor query with a threshold $w \leq \min\{\ell_{p'}, w_{i_k}\}$ can be answered by first counting the number h of 1s in the subarray $\hat{a}[0..\lfloor(w - w_{i_1})/\lfloor \log n \rfloor - 1]$ and, then, finding in $\mathcal{O}(1)$ time the predecessor of w in the fusion heap referred by P_{h+1} .

The array \hat{a} occupies $\mathcal{O}(1 + c/\log^2 n)$ space, which is $\mathcal{O}(1 + \ell_{p'}/\log^2 n)$ since $c \leq \ell_{p'}$. The computation of \hat{a} , which takes $\mathcal{O}(k + \ell_{p'}/\log^2 n)$ time, is the most time and space consuming part of the described construction; all other structures take $\mathcal{O}(k)$ time and space. We are to show that the computation of \hat{a} can sometimes be avoided if (almost) the same array was already constructed for a different path. The following lemma is the key for this optimization.

► **Lemma 8.** *Given a tree $I_{p'}$ for an irreducible position p' , consider its node x associated with a non-irreducible position $p > p'$. Let $s[r..n-1]$ be the closest irreducible lexicographical neighbour of $s[p..n-1]$ and let $c_p = \text{lcp}(r, p)$. Then, the subtree of $I_{p'}$ rooted at x coincides with the tree I_r in which all children of the root with weights $\geq c_p$ are cut, then all weights are increased by $p - p'$, and the weight of the root is set to the weight of x (see Figure 5).*


Proof. Denote by I the subtree rooted at x . Observe that all nodes of I are exactly all nodes of $I_{p'}$ associated with positions from a range $[p..p+i]$, for some $i < c_p$ (see Figure 5). The position $p + i + 1$ is either irreducible or $c_{p+i+1} + i + 1 \geq c_p$. In order to prove the lemma, it suffices to show that (1) all positions from $[r+1..r+i]$ are not irreducible, (2) $c_{p+j} = c_{r+j}$, for any $j \in [1..i]$, and (3) the position $r + i + 1$ is either irreducible or $c_{r+i+1} + i + 1 \geq c_p$.

(1) Suppose, to the contrary, that a position $r + j$, for some $j \in [1..i]$, is irreducible. Then, since $\text{lcp}(r, p) = c_p$ and $j \leq i < c_p$, we have $\text{lcp}(r + j, p + j) = c_p - j$. Therefore, c_{p+j} , the length of the common prefix of $s[p+j..n-1]$ and its closest irreducible lexicographical neighbour, must be at least $c_p - j$ (since it was assumed that $r + j$ is irreducible). But then the weight w_{p+j} of the node associated with $p + j$ is at least $(c_p - j) + (p + j - p') = c_p + p - p'$, which is equal to the weight $w_p = c_p + p - p'$ of x . Thus, x cannot be an ancestor of the node, which is a contradiction.

(2) For $j \in [1..i]$, $s[p+j..n-1]$ and $s[r+j..n-1]$ share a common prefix of length $c_p - j$ since $\text{lcp}(r, p) = c_p$ and $j \leq i < c_p$. Further, $c_{p+j} < c_p - j$ since the weight $w_{p+j} = c_{p+j} + p + j - p'$ is smaller than the weight $w_p = c_p + p - p'$ of its ancestor x . Suppose, without loss of generality, that $s[r+j..n-1]$ is lexicographically smaller than $s[p+j..n-1]$ (the case when it is greater is analogous). Then, neither of the suffixes $s[t..n-1]$ lexicographically lying between $s[r+j..n-1]$ and $s[p+j..n-1]$ can start with an irreducible position, for otherwise $c_{p+j} \geq \text{lcp}(t, p + j) \geq \text{lcp}(r + j, p + j) \geq c_p - j$. Therefore, the closest irreducible lexicographical neighbours of $s[p+j..n-1]$ and $s[r+j..n-1]$ coincide and $c_{p+j} = c_{r+j}$ ($< c_p - j \leq \text{lcp}(p + j, r + j)$).

(3) Suppose that $r + i + 1$ is not irreducible and, by contradiction, $c_{r+i+1} < c_p - i - 1$. The suffixes $s[p+i+1..n-1]$ and $s[r+i+1..n-1]$ have a common prefix of length $c_p - i - 1$ (note that $i + 1 \leq c_p$). The position $p + i + 1$ is not irreducible since otherwise $c_{r+i+1} \geq \text{lcp}(p + i + 1, r + i + 1) \geq c_p - i - 1$. Then, we have $c_{p+i+1} \geq c_p - i - 1$ because otherwise the node corresponding to $p + i + 1$ would have weight $w_{p+i+1} = c_{p+i+1} + p + i + 1 - p'$ that is smaller than the weight $w_p = c_p + p - p'$ of x and, thus, would be a descendant of x . Let $s[t..n-1]$ be the closest irreducible lexicographical neighbour of $s[p+i+1..n-1]$ so that $\text{lcp}(t, p + i + 1) = c_{p+i+1}$. Since $c_{p+i+1} \geq c_p - i - 1$ and $\text{lcp}(p + i + 1, r + i + 1) \geq c_p - i - 1$, we obtain $\text{lcp}(t, r + i + 1) \geq c_p - i - 1$. Because t is irreducible, we deduce that $c_{r+i+1} \geq c_p - i - 1$, which is a contradiction. ◀

An array \hat{a} corresponding to a heavy path in a tree $I_{p'}$, for an irreducible position p' , occupies $\mathcal{O}(1 + \ell_{p'}/\log^2 n)$ space. Recall that $\ell_{p'} = \text{plcp}[p']$ is an irreducible LCP value since p' is an irreducible position. Therefore, we can afford, for each tree $I_{p'}$, the construction and storage of the array \hat{a} corresponding to the unique heavy path containing the root of $I_{p'}$.



■ **Figure 5** For irreducible positions p' and r , the subtree of $I_{p'}$ highlighted by the left blue rectangle is isomorphic to I_r after cutting the children of the root of $I_{p'}$ outside of the right blue rectangle.

The overall space required for this is $\mathcal{O}(n + \sum_{p'} \ell_{p'} / \log^2 n)$, where the sum is through all irreducible positions p' , which is upperbounded by $\mathcal{O}(n + n / \log n) = \mathcal{O}(n)$ due to Lemma 2. Other heavy paths are processed as follows.

Suppose that we are to preprocess a heavy path $x_1 \rightarrow x_2 \rightarrow \dots \rightarrow x_k$ with node weights $w_{i_1}, w_{i_2}, \dots, w_{i_k}$ in a tree $I_{p'}$ such that x_k (the node closest to the root) is not the root of $I_{p'}$. We compute sets S_d and pointers P_h corresponding to the path in $\mathcal{O}(k)$ time. As for the array \hat{a} , we either construct it for the path from scratch or reuse a suitable array from another path; the details follow. Let $w_{i_m} = \max\{w_{i_j} : w_{i_j} \leq \ell_{p'}\}$. As was discussed, the array \hat{a} encodes, in a sense, a predecessor data structure for the weights $w_{i_1}, w_{i_2}, \dots, w_{i_m}$. In order to have some flexibility for the reuse of arrays, we modify this scheme slightly so that sometimes \hat{a} does not encode the last two weights $w_{i_{m-1}}$ and w_{i_m} . The algorithm that answers a predecessor query for a threshold $w \leq \min\{\ell_{p'}, w_{i_k}\}$ on the path is altered accordingly: we first compare w to $w_{i_{m-1}}$ and w_{i_m} , and only then, if necessary, use the array \hat{a} as described above.

Let x_k correspond to a non-irreducible position p and let $s[r..n-1]$ be the closest irreducible lexicographical neighbour of $s[p..n-1]$. As was shown in Section 4, the position r can be found in $\mathcal{O}(1)$ time. By Lemma 8, the subtree I rooted at x_k is isomorphic to the tree I_r in which all children of the root with weights greater than or equal to c_p are cut, then all weights are increased by $p - p'$, and the root weight is set to w_{i_k} . Denote this isomorphism by ϕ . Note that $\phi(x_k)$ is the root of I_r . The main corollary of Lemma 8 is that in the subtree I any edge $u \rightarrow v$ that is not incident to x_k is heavy iff the corresponding edge $\phi(u) \rightarrow \phi(v)$ in I_r is heavy. Therefore, all edges in the path $\phi(x_1) \rightarrow \phi(x_2) \rightarrow \dots \rightarrow \phi(x_k)$ are heavy, except, possibly, the last edge $\phi(x_{k-1}) \rightarrow \phi(x_k)$, and no heavy edge enters the node $\phi(x_1)$ in I_r . By Lemma 8, the weights of the nodes $\phi(x_1), \phi(x_2), \dots, \phi(x_{k-1})$ are $w_{i_1} - (p - p')$, $w_{i_2} - (p - p')$, \dots , $w_{i_{k-1}} - (p - p')$, respectively. We proceed further in three separate cases.

Heavy $\phi(x_{k-1}) \rightarrow \phi(x_k)$ and ℓ_r is large enough. Suppose that the edge $\phi(x_{k-1}) \rightarrow \phi(x_k)$ is heavy. Then, $\phi(x_1) \rightarrow \phi(x_2) \rightarrow \dots \rightarrow \phi(x_k)$ is the unique heavy path of I_r that contains the root. Let $\hat{a}_\phi[0..c_\phi]$ be an array for predecessor queries that was explicitly stored in $\mathcal{O}(1 + \ell_r / \log^2 n)$ space for the path $\phi(x_1) \rightarrow \phi(x_2) \rightarrow \dots \rightarrow \phi(x_k)$. By definition, for any $d \in [0..c_\phi]$, we have $\hat{a}_\phi[d] = 1$ iff, for some $j \in [1..k]$, the number $(w_{i_j} - (p - p')) - (w_{i_1} - (p - p')) = w_{i_j} - w_{i_1}$ lies in the range $[d \lceil \log n \rceil .. (d + 1) \lceil \log n \rceil - 1]$. Since the latter is also a criterium for $\hat{a}[d] = 1$, the array \hat{a}_ϕ can therefore be reused to imitate \hat{a} if \hat{a}_ϕ is sufficiently long to “encode” the weights $w_{i_1}, w_{i_2}, \dots, w_{i_{m-2}}$. More precisely, this is the case iff $\ell_r \geq w_{i_{m-2}} - (p - p')$. Thus, if the edge $\phi(x_{k-1}) \rightarrow \phi(x_k)$ is heavy and $\ell_r \geq w_{i_{m-2}} - (p - p')$, then the overall preprocessing of the path $x_1 \rightarrow x_2 \rightarrow \dots \rightarrow x_k$ takes only $\mathcal{O}(k)$ time since we can store a pointer to the array \hat{a}_ϕ and reuse \hat{a}_ϕ to imitate the array \hat{a} .

Unfortunately, neither of these two conditions necessarily holds in general: the edge $\phi(x_{k-1}) \rightarrow \phi(x_k)$ might be light in I_r and ℓ_r might be less than $w_{i_{m-2}} - (p - p')$.

Light $\phi(x_{k-1}) \rightarrow \phi(x_k)$ and ℓ_r is large enough. Suppose that the edge $\phi(x_{k-1}) \rightarrow \phi(x_k)$ is light in I_r and $\ell_r \geq w_{i_{m-2}} - (p - p')$. By Lemma 8, the edge necessarily becomes heavy if we cut all children of the root of I_r with weights greater than or equal to c_p . We assign numbers $1, 2, \dots$ to the children of each node in the trees I_r and $I_{p'}$ according to the order in which they were attached during the construction of the trees, i.e., in the increasing order of their associated positions. It follows from the proof of Lemma 8 that if x_{k-1} is the h th child of x_k , then $\phi(x_{k-1})$ is the h th child of the root of I_r . Therefore, the node $\phi(x_{k-1})$ can be located in $\mathcal{O}(1)$ time.

We associate with each child of the root of I_r a pointer, initially set to null. If the pointer in $\phi(x_{k-1})$ is still null at the time we access it while preprocessing the heavy path $x_1 \rightarrow x_2 \rightarrow \dots \rightarrow x_k$, then we create from scratch a new array \hat{a}_ϕ for the heavy path $\phi(x_1) \rightarrow \phi(x_2) \rightarrow \dots \rightarrow \phi(x_{k-1})$ in $\mathcal{O}(k + \ell_r / \log^2 n)$ time and set the pointer of $\phi(x_{k-1})$ to refer to this array. Since $\ell_r \geq w_{i_{m-2}} - (p - p')$, the array \hat{a}_ϕ can be reused to imitate \hat{a} for the path $x_1 \rightarrow x_2 \rightarrow \dots \rightarrow x_k$ in the same way as was described above. If the pointer in the node $\phi(x_{k-1})$ is not null, then it already refers to a suitable array \hat{a}_ϕ that can be reused (note that $\phi(x_1) \rightarrow \phi(x_2) \rightarrow \dots \rightarrow \phi(x_{k-1})$ is the unique heavy path of the tree I_r that contains the node $\phi(x_{k-1})$). Thus, the preprocessing takes $\mathcal{O}(k)$ time plus, if necessary, $\mathcal{O}(k + \ell_r / \log^2 n)$ time required to construct a new array \hat{a}_ϕ .

The following lemma shows that a non-null pointer to an array \hat{a}_ϕ might be assigned to at most $\log n$ distinct children of the root of I_r (namely, those children that become connected to the root by a heavy edge after a number of children with greater weights were removed). This implies that the overall construction time for all the arrays \hat{a}_ϕ throughout the whole algorithm is $\mathcal{O}(n + \sum_r (\ell_r / \log^2 n) \log n) = \mathcal{O}(n)$, where the sum is through all irreducible positions r (so that $\sum_r \ell_r = \mathcal{O}(n \log n)$ by Lemma 2).

► **Lemma 9.** *Let S be a tree with at most n nodes whose root r has m children ordered arbitrarily. Suppose that we remove the children of the root (with the subtrees rooted at them) from right to left, one by one, thus producing trees S_1, S_2, \dots, S_m . Let z_i be a node of the tree S_i such that z_i is connected to r by a heavy edge, or z_i is r itself if r has no incident heavy edges in S_i . Then, the set $\{z_1, z_2, \dots, z_m\}$ contains at most $\log n$ distinct nodes.*

Proof. If $z_i \rightarrow r$ is a heavy edge in S_i , then $\text{size}(z_i) > \text{size}(r)/2$ in S_i . Therefore, if we cut any other child of r in S_i , only the number $\text{size}(r)$ decreases and, thus, the edge remains heavy. But when we remove z_i and its subtree from S_i , the number $\text{size}(r)$ decreases by more than half. Such halving may happen at most $\log n - 1$ times, and the result follows. ◀

Small ℓ_r . Suppose that $\ell_r < w_{i_{m-2}} - (p - p')$. Let \bar{p} be the position associated with the node x_{k-1} in the tree $I_{p'}$ and let $s[\bar{p}..n-1]$ be the closest irreducible lexicographical neighbour of $s[\bar{p}..n-1]$. By analogy to the isomorphism ϕ , we define using Lemma 8 an isomorphism ψ that maps the subtree of $I_{p'}$ rooted at the node x_{k-1} onto a “pruned” tree $I_{\bar{r}}$ in which all children of the root with weights greater than or equal to $c_{\bar{p}}$ are cut. Note that $\psi(x_{k-1})$ is the root of $I_{\bar{r}}$. By Lemma 8, the weights of the nodes $\psi(x_1), \psi(x_2), \dots, \psi(x_{k-2})$ are $w_{i_1} - (\bar{p} - p')$, $w_{i_2} - (\bar{p} - p')$, \dots , $w_{i_{k-2}} - (\bar{p} - p')$, respectively.

It turns out that, instead of imitating the array \hat{a} for the path $x_1 \rightarrow x_2 \rightarrow \dots \rightarrow x_k$ using an array \hat{a}_ϕ from the tree I_r , one can in quite the same way imitate \hat{a} using an appropriate array \hat{a}_ψ from $I_{\bar{r}}$, which must be sufficiently long in the case $\ell_r < w_{i_{m-2}} - (p - p')$. More precisely, we are to prove that $\ell_{\bar{r}} \geq w_{i_{m-2}} - (\bar{p} - p')$, which guarantees that essentially the same approach works well: if the edge $\psi(x_{k-2}) \rightarrow \psi(x_{k-1})$ is heavy, then \hat{a}_ψ is an array associated with the unique heavy path containing the root of $I_{\bar{r}}$ and, since $\ell_{\bar{r}} \geq w_{i_{m-2}} - (\bar{p} - p')$, the

array \hat{a}_ψ is long enough to imitate \hat{a} ; if the edge $\phi(x_{k-2}) \rightarrow \psi(x_{k-1})$ is light, then we either reuse a suitable array \hat{a}_ψ that was already stored in the child $\psi(x_{k-2})$ of the root of $I_{\bar{r}}$ (again \hat{a}_ψ can imitate \hat{a} since $\ell_{\bar{r}} \geq w_{i_{m-2}} - (\bar{p} - p')$) or we create this array \hat{a}_ψ from scratch for the path $\psi(x_1) \rightarrow \psi(x_2) \rightarrow \dots \rightarrow \psi(x_{k-2})$ in $\mathcal{O}(k + \ell_{\bar{r}}/\log^2 n)$ time and store a pointer to it in $\psi(x_{k-2})$. We do not go further into details since they are essentially the same as above.

We actually prove a stronger condition $\ell_{\bar{r}} \geq w_{i_{m-1}} - (\bar{p} - p')$, which implies $\ell_{\bar{r}} \geq w_{i_{m-2}} - (\bar{p} - p')$ since $w_{i_{m-1}} > w_{i_{m-2}}$. Recall that $w_{i_{m-1}} = c_{\bar{p}} + \bar{p} - p'$. Hence, the stronger condition is equivalent to $\ell_{\bar{r}} \geq c_{\bar{p}}$. If the suffix $s[\bar{r}..n-1]$ is lexicographically larger than $s[\bar{p}..n-1]$, then we immediately obtain $\ell_{\bar{r}} = \text{plcp}[\bar{r}] \geq \text{lcp}(\bar{p}, \bar{r}) = c_{\bar{p}}$. Assume that $s[\bar{r}..n-1]$ is lexicographically smaller than $s[\bar{p}..n-1]$. Let us show that this case is actually impossible since it contradicts the condition of “small ℓ_r ” that was assumed in the beginning: $\ell_r < w_{i_{m-2}} - (p - p')$. Denote $j = \bar{p} - p$. Since $\text{lcp}(p, r) = c_p$ and $j < c_p$, we obtain $\text{lcp}(\bar{p}, r+j) = c_p - j$. Further, $c_{\bar{p}} < c_p - j$ because the weight $w_{i_{m-1}} = c_{\bar{p}} + \bar{p} - p'$ of the node x_{m-1} is less than the weight $w_{i_k} = c_p + p - p'$ of x_k . Since $\text{lcp}(\bar{p}, \bar{r}) = c_{\bar{p}} < c_p - j = \text{lcp}(\bar{p}, r+j)$, we obtain $\text{lcp}(\bar{r}, r+j) = c_{\bar{p}} < \text{lcp}(\bar{p}, r+j)$ and, further, since we assumed that $s[\bar{r}..n-1]$ is lexicographically smaller than $s[\bar{p}..n-1]$, the suffix $s[\bar{r}..n-1]$ is lexicographically smaller than $s[r+j..n-1]$. Hence, $\ell_{r+j} \geq \text{lcp}(\bar{r}, r+j) = c_{\bar{p}}$. It follows from Lemma 4 that $\ell_r = \ell_{r+j} + j$. Then, $\ell_r \geq c_{\bar{p}} + j$. Conversely, we deduce $\ell_r < w_{i_{m-2}} - (p - p') < w_{i_{m-1}} - (p - p')$ and, substituting $w_{i_{m-1}} = c_{\bar{p}} + p + j - p'$, we obtain $\ell_r < c_{\bar{p}} + j$, which is a contradiction.

To sum up, we spend linear time preprocessing each heavy path in every tree $I_{p'}$ plus $\mathcal{O}(n)$ total time to construct arrays \hat{a}_ϕ and \hat{a}_ψ , each of which is associated with one of $\log n$ particular children of the root in one of the trees (according to Lemma 9). Therefore, since all the trees $I_{p'}$ contain n nodes in total, the overall time is $\mathcal{O}(n)$.

6 Concluding Remarks

We believe that the presented solution, while certainly still quite complicated and impractical, is more implementable than that of [13]. We expect that, with additional combinatorial insights, one can simplify it even further, perhaps eventually arriving at a practical data structure for the problem. A number of natural and less vaguely formulated problems also arise. For example, can we maintain the weighted ancestor data structure during an online construction of the suffix tree? Is it possible to reduce the space usage and support weighted ancestors in compact suffix trees [24, 27] with $\mathcal{O}(n)$ or $\mathcal{O}(n \log \sigma)$ additional *bits* of space, where σ is the alphabet size?

To the best of our knowledge, the lemma about the sum of irreducible LCP values has found relatively few applications other than in the construction of LCP and PLCP arrays (e.g., see references in [16]): [2, 3, 18, 29] (the application in [18], however, is somewhat spectacular). The techniques developed in our paper might be interesting by themselves and pave the way to more applications of irreducible LCPs in algorithms on strings.

References

- 1 A. Amir, G.M. Landau, M. Lewenstein, and D. Sokol. Dynamic text and static pattern matching. *ACM Transactions on Algorithms (TALG)*, 3(2):19–31, 2007. doi:10.1145/1240233.1240242.
- 2 G. Badkobeh, P. Gawrychowski, J. Kärkkäinen, S. J. Puglisi, and B. Zhukova. Tight upper and lower bounds on suffix tree breadth. *Theoretical Computer Science*, 854:63–67, 2021. doi:10.1016/j.tcs.2020.11.037.

- 3 G. Badkobeh, J. Kärkkäinen, S.J. Puglisi, and B. Zhukova. On suffix tree breadth. In *Proc. 24th Symposium on String Processing and Information Retrieval (SPIRE)*, volume 10508 of *LNCS*, pages 68–73. Springer, 2017. doi:10.1007/978-3-319-67428-5_6.
- 4 M. A. Bender and M. Farach-Colton. The LCA problem revisited. In *4th Latin American Symposium on Theoretical Informatics (LATIN)*, pages 88–94. Springer, 2000. doi:10.1007/10719839_9.
- 5 M.A. Bender and M. Farach-Colton. The level ancestor problem simplified. *Theoretical Computer Science*, 321(1):5–12, 2004. doi:10.1016/j.tcs.2003.05.002.
- 6 O. Berkman and U. Vishkin. Finding level-ancestors in trees. *Journal of Computer and System Sciences*, 48(2):214–230, 1994. doi:10.1016/S0022-0000(05)80002-9.
- 7 S. Biswas, A. Ganguly, R. Shah, and S.V. Thankachan. Ranked document retrieval with forbidden pattern. In *Proc. 26th Annual Symposium on Combinatorial Pattern Matching (CPM)*, pages 77–88. Springer, 2015. doi:10.1007/978-3-319-19929-0_7.
- 8 S. Biswas, A. Ganguly, R. Shah, and S.V. Thankachan. Ranked document retrieval for multiple patterns. *Theoretical Computer Science*, 746:98–111, 2018. doi:10.1016/j.tcs.2018.06.029.
- 9 M. Burrows and D. J. Wheeler. A block-sorting lossless data compression algorithm. Technical Report 124, Digital Equipment Corporation, Palo Alto, California, 1994.
- 10 David Clark. *Compact pat trees*. PhD thesis, University of Waterloo, 1997.
- 11 M. Farach and S. Muthukrishnan. Perfect hashing for strings: Formalization and algorithms. In *7th Annual Symposium on Combinatorial Pattern Matching (CPM)*, pages 130–140. Springer, 1996. doi:10.1007/3-540-61258-0_11.
- 12 M. L. Fredman and D. E. Willard. Surpassing the information theoretic bound with fusion trees. *Journal of Computer and System Sciences*, 47(3):424–436, 1993. doi:10.1016/0022-0000(93)90040-4.
- 13 P. Gawrychowski, M. Lewenstein, and P.K. Nicholson. Weighted ancestors in suffix trees. In *Proc. 22nd Annual European Symposium on Algorithms (ESA)*, pages 455–466. Springer, 2014. doi:10.1007/978-3-662-44777-2_38.
- 14 D. Harel and R. E. Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM Journal on Computing*, 13(2):338–355, 1984. doi:10.1137/0213024.
- 15 G. Jacobson. Space-efficient static trees and graphs. In *Proc. 30th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 549–554. IEEE, 1989. doi:10.1109/SFCS.1989.63533.
- 16 J. Kärkkäinen, D. Kempa, and M. Piątkowski. Tighter bounds for the sum of irreducible LCP values. *Theoretical Computer Science*, 656:265–278, 2016. doi:10.1016/j.tcs.2015.12.009.
- 17 J. Kärkkäinen, G. Manzini, and S.J. Puglisi. Permuted longest-common-prefix array. In *Proc. 20th Annual Symposium on Combinatorial Pattern Matching (CPM)*, pages 181–192. Springer, 2009. doi:10.1007/978-3-642-02441-2_17.
- 18 D. Kempa and T. Kociumaka. Resolution of the Burrows–Wheeler transform conjecture. *arXiv preprint*, 2019. arXiv:1910.10631.
- 19 D. Kempa, A. Policriti, N. Prezza, and E. Rotenberg. String attractors: Verification and optimization. In *Proc. 26th Annual European Symposium on Algorithms (ESA)*, volume 112 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 52:1–52:13. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018. doi:10.4230/LIPIcs.ESA.2018.52.
- 20 T. Kociumaka, M. Kubica, J. Radoszewski, W. Rytter, and T. Waleń. A linear-time algorithm for seeds computation. *ACM Transactions on Algorithms (TALG)*, 16(2):1–23, 2020. doi:10.1145/3386369.
- 21 T. Kopelowitz, G. Kucherov, Y. Nekrich, and T. Starikovskaya. Cross-document pattern matching. *Journal of Discrete Algorithms*, 24:40–47, 2014. doi:10.1016/j.jda.2013.05.002.
- 22 T. Kopelowitz and M. Lewenstein. Dynamic weighted ancestors. In *Proc. 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 565–574. ACM-SIAM, 2007. doi:10.5555/1283383.1283444.

- 23 U. Manber and G. Myers. Suffix arrays: a new method for on-line string searches. *SIAM Journal on Computing*, 22(5):935–948, 1993. doi:10.1137/0222058.
- 24 J. I. Munro, G. Navarro, and Y. Nekrich. Space-efficient construction of compressed indexes in deterministic linear time. In *Proc. 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 408–424. SIAM, 2017. doi:10.1137/1.9781611974782.26.
- 25 M. Pătraşcu and M. Thorup. Time-space trade-offs for predecessor search. In *Proc. 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 232–240. ACM, 2006. doi:10.1145/1132516.1132551.
- 26 M. Pătraşcu and M. Thorup. Dynamic integer sets with optimal rank, select, and predecessor search. In *Proc. 55th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 166–175. IEEE, 2014. doi:10.1109/FOCS.2014.26.
- 27 K. Sadakane. Compressed suffix trees with full functionality. *Theory of Computing Systems*, 41(4):589–607, 2007. doi:10.1007/s00224-006-1198-x.
- 28 D. D. Sleator and R. E. Tarjan. A data structure for dynamic trees. *Journal of Computer and System Sciences*, 26(3):362–391, 1983. doi:10.1016/0022-0000(83)90006-5.
- 29 N. Välimäki and S.J. Puglisi. Distributed string mining for high-throughput sequencing data. In *Proc. 12th International Workshop on Algorithms in Bioinformatics (WABI)*, LNCS 7534, pages 441–452. Springer, 2012. doi:10.1007/978-3-642-33122-0_35.
- 30 P. van Emde Boas, R. Kaas, and E. Zijlstra. Design and implementation of an efficient priority queue. *Mathematical systems theory*, 10(1):99–127, 1976. doi:10.1007/BF01683268.
- 31 D. E. Willard. Log-logarithmic worst-case range queries are possible in space $\Theta(N)$. *Information Processing Letters*, 17(2):81–84, 1983. doi:10.1016/0020-0190(83)90075-3.