

Longest Common Extensions in Trees

Philip Bille^{*1}, Paweł Gawrychowski^{**2}, Inge Li Gørtz^{*1},
Gad M. Landau^{***3,4}, and Oren Weimann^{†3}

¹ DTU Informatics, Denmark. phbi@dtu.dk, inge@dtu.dk

² University of Warsaw, Poland. gawry@mimuw.edu.pl

³ University of Haifa, Israel, landau@cs.haifa.ac.il, oren@cs.haifa.ac.il

⁴ NYU, USA.

Abstract. The longest common extension (LCE) of two indices in a string is the length of the longest identical substrings starting at these two indices. The LCE problem asks to preprocess a string into a compact data structure that supports fast LCE queries.

In this paper we generalize the LCE problem to trees and suggest a few applications of LCE in trees to tries and XML databases. Given a labeled and rooted tree T of size n , the goal is to preprocess T into a compact data structure that support the following LCE queries between subpaths and subtrees in T . Let v_1, v_2, w_1 , and w_2 be nodes of T such that w_1 and w_2 are descendants of v_1 and v_2 respectively.

- $\text{LCE}_{PP}(v_1, w_1, v_2, w_2)$: (path-path LCE) return the longest common prefix of the paths $v_1 \rightsquigarrow w_1$ and $v_2 \rightsquigarrow w_2$.
- $\text{LCE}_{PT}(v_1, w_1, v_2)$: (path-tree LCE) return maximal path-path LCE of the path $v_1 \rightsquigarrow w_1$ and any path from v_2 to a descendant leaf.
- $\text{LCE}_{TT}(v_1, v_2)$: (tree-tree LCE) return a maximal path-path LCE of any pair of paths from v_1 and v_2 to descendant leaves.

We present the first non-trivial bounds for supporting these queries. For LCE_{PP} queries, we present a linear-space solution with $O(\log^* n)$ query time. For LCE_{PT} queries, we present a linear-space solution with $O((\log \log n)^2)$ query time, and complement this with a lower bound showing that any path-tree LCE structure of size $O(n \text{polylog}(n))$ must necessarily use $\Omega(\log \log n)$ time to answer queries. For LCE_{TT} queries, we present a time-space trade-off, that given any parameter $\tau, 1 \leq \tau \leq n$, leads to an $O(n\tau)$ space and $O(n/\tau)$ query-time solution. This is complemented with a reduction to the the set intersection problem implying that a fast linear space solution is not likely to exist.

1 Introduction

Given a string S , the *longest common extension* (LCE) of two indices is the length of the longest identical substring starting at these indices. The *longest*

* Partially supported by the Danish Research Council (DFR – 4005-00267, DFR – 1323-00178) and the Advanced Technology Foundation.

** Work done while the author held a post-doctoral position at the Warsaw Center of Mathematics and Computer Science.

*** Partially supported by the National Science Foundation Award 0904246, Israel Science Foundation grant 571/14, Grant No. 2008217 from the United States-Israel Binational Science Foundation (BSF) and DFG.

† Partially supported by the Israel Science Foundation grant 794/13.

common extension problem (LCE problem) is to preprocess S into a compact data structure supporting fast LCE queries. The LCE problem is a basic primitive in a wide range of string matching problems such as approximate string matching, finding exact and approximate tandem repeats, and finding palindromes [2, 9, 15, 18–20]. The classic textbook solution to the LCE problem on strings combines a suffix tree with a nearest common ancestor (NCA) data structure leading to a linear space and constant query-time solution [14].

In this paper we study generalizations of the LCE problem to trees. The goal is to preprocess an edge-labeled, rooted tree T to support the various LCE queries between paths in T . Here a path starts at a node v and ends at a descendant of v , and the LCEs are on the strings obtained by concatenating the characters on the edges of the path from top to bottom (each edge contains a single character). We consider path-path LCE queries between two specified paths in T , path-tree LCE queries defined as the maximal path-path LCE of a given path and *any* path starting at a given node, and tree-tree LCE queries defined as the maximal path-path LCE between *any* pair of paths starting from two given nodes. We next define these problems formally.

Tree LCE Problems. Let T be an edge-labeled, rooted tree with n nodes. We denote the subtree rooted at a node v by $T(v)$, and given nodes v and w such that w is in $T(v)$ the path going down from v to w is denoted $v \rightsquigarrow w$. A *path prefix* of $v \rightsquigarrow w$ is any subpath $v \rightsquigarrow u$ such that u is on the path $v \rightsquigarrow w$. Two paths $v_1 \rightsquigarrow w_1$ and $v_2 \rightsquigarrow w_2$ *match* if concatenating the labels of all edges in the paths gives the same string. Given nodes v_1, w_1 such that $w_1 \in T(v_1)$ and nodes v_2, w_2 such that $w_2 \in T(v_2)$ define the following queries:

- $\text{LCE}_{PP}(v_1, w_1, v_2, w_2)$: (path-path LCE) return the longest common matching prefix of the paths $v_1 \rightsquigarrow w_1$ and $v_2 \rightsquigarrow w_2$.
- $\text{LCE}_{PT}(v_1, w_1, v_2)$: (path-tree LCE) return the maximal path-path LCE of the path $v_1 \rightsquigarrow w_1$ and any path from v_2 to a descendant leaf.
- $\text{LCE}_{TT}(v_1, v_2)$: (tree-tree LCE) return a maximal path-path LCE of any pair of paths from v_1 and v_2 to descendant leaves.

We assume that the output of the queries is reported compactly as the endpoint(s) of the LCE. This allows us to report the shared path in constant time per edge. Furthermore, we will assume w.l.o.g. that for each node v in T , all the edge-labels to children of v are distinct. If this is not the case, then we can merge all identical edges of a node to its children in linear time, without affecting the result of all the above LCE queries.

We note that the direction of the paths in T is important for the LCE queries. In the above LCE queries, the paths start from a node and go downwards. If we instead consider paths from a node going upwards towards the root of T , the problem is easier and can be solved in linear space and constant query-time by combining Breslauer’s suffix tree of a tree [7] with a nearest common ancestor (NCA) data structure [16].

Our Results. First consider the LCE_{PP} and LCE_{PT} problems. To answer an $\text{LCE}_{PP}(v_1, w_1, v_2, w_2)$ query, a straightforward solution is to traverse both paths in parallel-top down. Similarly, to answer an $\text{LCE}_{PT}(v_1, w_1, v_2)$ query we can

traverse $v_1 \rightsquigarrow w_1$ top-down while traversing the matching path from v_2 (recall that all edges to a child are distinct and hence the longest matching path is unique). This approach leads to a linear-space solution with $O(h)$ query-time to both problems, where h is the height of T . Note that for worst-case trees we have that $h = \Omega(n)$.

We show the following results. For LCE_{PP} we give a linear $O(n)$ space and $O(\log^* n)$ query-time solution. For LCE_{PT} we give a linear $O(n)$ space and $O((\log \log n)^2)$ query-time solution, and complement this with a lower bound stating that any LCE_{PT} solution using $O(n \text{ polylog}(n))$ space must necessarily have $\Omega(\log \log n)$ query time.

Next consider the LCE_{TT} problem. Here, the simple top down traversal does not work and it seems that substantially different ideas are needed. We first show a reduction from the *set-intersection problem*, i.e., preprocessing a family of sets of total size n to support disjointness queries between any pairs of sets. In particular, the reduction implies that a fast linear space solution is not likely assuming a widely believed conjecture on the complexity of the set-intersection problem. We complement this result with a time-space trade-off that achieves $O(n\tau)$ space and $O(n/\tau)$ query time for any parameter $1 \leq \tau \leq n$.

All results assume the standard word RAM model with word size $\Theta(\log n)$. We also assume the alphabet is either sorted or is linear-time sortable.

Applications. We suggest a few immediate applications of LCE in trees. Consider a set of strings $\mathcal{S} = \{S_1, \dots, S_k\}$ of total length $\sum_{i=1}^k |S_i| = N$ and let T be the *trie* of \mathcal{S} of size n , i.e., T is the labeled, rooted tree obtained by merging shared prefixes in \mathcal{S} maximally. If we want to support LCE queries between suffixes of strings in \mathcal{S} , the standard approach is to build a generalized suffix tree for the strings and combine it with an NCA data structure. This leads to a solution using $O(N)$ space and $O(1)$ query time. We can instead implement the LCE query between the suffixes of strings in \mathcal{S} as an LCE_{PP} on the trie T . With our data structure for LCE_{PP} , this leads to a solution using $O(n)$ space and $O(\log^* n)$ query time. In general, n can be significantly smaller than N , depending on the amount of shared prefixes in \mathcal{S} . Hence, this solution provides a more space-efficient representation of \mathcal{S} at the expense of a tiny increase in query time. An LCE_{PT} query on T corresponds to computing a maximal LCE of a suffix of a string in \mathcal{S} with suffixes of strings in \mathcal{S} sharing a common prefix. An LCE_{TT} query on T corresponds to computing a maximal LCE over pairs of suffixes of strings in \mathcal{S} that share a common prefix. To the best of our knowledge these queries are novel one-to-many and many-to-many LCE queries. Since tries are a basic data structure for storing strings we expect these queries to be of interest in a number of applications.

Another interesting application is using LCE in trees as a query primitive for XML data. XML documents can be viewed as a labeled tree and typical queries (e.g., XPath queries) involve traversing and identifying paths in the tree. The LCE queries provide simple and natural primitives for comparing paths and subtrees without explicit traversal. For instance, our solution for LCE_{PT} queries can be used to quickly identify the “best match” of a given path in a subtree.

2 Preliminaries

Given a node v and an integer $d \geq 0$, the *level ancestor* of v at depth d , denoted $LA(v, d)$ is the ancestor of v at depth d . We explicitly compute and store the depth of every node v , denoted $\text{depth}(v)$. Given a pair of nodes v and w the *nearest common ancestor* of v and w , denoted $NCA(v, w)$, is the common ancestor of v and w of greatest depth. Both LA and NCA queries can be supported in constant time with a linear space data structures, see e.g., [1, 4–6, 10, 11, 13, 16]

Finally, the *suffix tree of a tree* [7, 17, 24] is the compressed trie of all suffixes of leaf-to-root paths in T . The suffix tree of a tree uses $O(n)$ space and can be constructed in $O(n \log \log n)$ time for general alphabets [24]. Note that the suffix tree of a tree combined with NCA can support LCE queries in constant time for paths going upwards. Since we consider paths going downwards, we will only use the suffix tree to check (in constant time) if two paths are completely identical.

We also need the following three primitives. *Range minimum queries*: A list of n numbers a_1, a_2, \dots, a_n can be augmented with $2n + o(n)$ bits of additional data in $O(n)$ time, so that for any $i \leq j$ the position of the smallest number among a_i, a_{i+1}, \dots, a_j can be found in $O(1)$ time [11]. *Predecessor queries*: Given a sorted collection of n integers from $[0, U)$, a structure of size $O(n)$ answering predecessor queries in $O(\log \log U)$ time can be constructed in time $O(n)$ [25], where a predecessor query locates, for a given x , the largest $y \leq x$ such that $y \in S$. Finally, *Perfect hashing*: given a collection S of n integers a perfect hash table can be constructed in expected $O(n)$ time [12], where a perfect hash table checks, for a given x , if $x \in S$, and if so returns its associated data in $O(1)$ time. The last result can be made deterministic at the expense of increasing the preprocessing time to $O(n \log \log n)$ [23], but then we need one additional step in our solution for the path-tree LCE as to ensure $O(n)$ total construction time.

3 Difference Covers for Trees

In this section we introduce a generalization of difference covers from strings to trees. This will be used to decrease the space of our data structures. We believe it is of independent interest.

Lemma 1. *For any tree T with n nodes and a parameter x , it is possible to mark $2n/x$ nodes of T , so that for any two nodes $u, v \in T$ at (possibly different) depths at least x^2 , there exists $d \leq x^2$ such that the d -th ancestors of both u and v are marked. Furthermore, such d can be calculated in $O(1)$ time and the set of marked nodes can be determined in $O(n)$ time.*

Proof. We distinguish between two types of marked nodes. Whether a node v is marked or not depends only on its depth. The marked nodes are determined as follows:

Type I. For every $i = 0, 1, \dots, x - 1$, let V_i be the set of nodes at depth leaving a remainder of i when divided by x . Because $\bigcup_i V_i = T$ and all V_i 's are disjoint, there exists $r_1 \in [0, x - 1]$ such that $|V_{r_1}| \leq n/x$. Then v is a type I marked node iff $\text{depth}(v) = r_1 \bmod x$.

Type II. For every $i = 0, 1, \dots, x - 1$, let V_i be the set of nodes v such that $\lfloor \text{depth}(v)/x \rfloor$ leaves a remainder of i when divided by x . By the same argument as above, there exists $r_2 \in [0, x - 1]$ such that $|V_{r_2}| \leq n/x$. Then v is a type II marked node iff $\lfloor \text{depth}(v)/x \rfloor = r_2 \pmod{x}$.

Now, given two nodes u and v at depths at least x^2 , we need to show that there exists an appropriate $d \leq x^2$. Let $\text{depth}(u) = t_1 \pmod{x}$ and choose $d_1 = t_1 + x - r_1$. Then the d_1 -th ancestor of u is a type I marked node, because its depth is $\text{depth}(u) - d_1 = \text{depth}(u) - (t_1 + x - r_1) = \text{depth}(u) - t_1 - x + r_1$, which leaves a remainder of r_1 when divided by x . Our d will be of the form $d_1 + d_2x$. Observe that regardless of the value of d_2 , we can be sure that the d -th ancestor of u is a type I marked node. Let v' be the d_1 -th ancestor of v , $\lfloor \text{depth}(v')/x \rfloor = t_2 \pmod{x}$ and choose $d_2 = t_2 + x - r_2$. The (d_2x) -th ancestor of v' is a type II marked node, because $\lfloor (\text{depth}(v') - d_2x)/x \rfloor = \lfloor \text{depth}(v')/x \rfloor - t_2 - x + r_2$, which leaves a remainder of r_2 when divided by x . Therefore, choosing $d = d_1 + d_2x$ guarantees that $d \leq x - 1 + x(x - 1) < x^2$, so the d -th ancestors of u and v are both defined, the d -th ancestor of u is a type I marked node, and the d -th ancestor of v is a type II marked node.

The total number of marked nodes is clearly at most $2n/x$, and the values of r and r' can be determined by a single traversal of T . To determine d , we only need to additionally know $\text{depth}(u)$ and $\text{depth}(v)$ and perform a few simple arithmetical operations. \square

Remark. Our difference cover has the following useful property: whether a node v is marked or not depends only on the value of $\text{depth}(v) \pmod{x^2}$. Hence, if a node at depth at least x^2 is marked then so is its (x^2) -th ancestor. Similarly, if a node is marked, so are all of its descendants at distance x^2 .

4 Path-Path LCE

In this section we prove the following theorem.

Theorem 1. *For a tree T with n nodes, a data structure of size $O(n)$ can be constructed in $O(n)$ time to answer path-path LCE queries in $O(\log^* n)$ time.*

We start with a simple preliminary $O(n \log n)$ -space $O(1)$ -query data structure which will serve as a starting point for the more complicated final implementation. We note that a data structure with similar guarantees to Lemma 2 is also implied from [3].

Lemma 2. *For a tree T with n nodes, a data structure of size $O(n \log n)$ can be constructed in $O(n \log n)$ time to answer path-path LCE queries in $O(1)$ time.*

Proof. The structure consists of $\log n$ separate parts, each of size $O(n)$. The k -th part answers in $O(1)$ time path-path LCE queries such that both paths are of the same length 2^k . This is enough to answer a general path-path LCE query in the same time complexity, because we can first truncate the longer path so that both paths are of the same length ℓ , then calculate k such that $2^k \leq \ell < 2^{k+1}$. Then we have two cases:

1. The prefixes of length 2^k of both paths are different. Then replacing the paths by their prefixes of length 2^k does not change the answer.
2. The prefixes of length 2^k of both paths are the same. Then replacing the paths by their suffixes of length 2^k does not change the answer.

We can check if the prefixes are the same and then (with level ancestor queries) reduce the query so that both paths are of the same length 2^k , all in $O(1)$ time.

Consider all paths of length 2^k in the tree. There are at most n of them, because every node u creates at most one new path $\text{LA}(v, \text{depth}(v) - 2^k) \rightsquigarrow v$. We lexicographically sort all such paths and store the longest common extension of every two neighbours on the sorted list. Additionally, we augment the longest common extensions with a range minimum query structure, and keep at every v the position of the path $\text{LA}(v, \text{depth}(v) - 2^k) \rightsquigarrow v$ (if any) on the sorted list. This allows us to answer $\text{LCE}_{PP}(\text{LA}(u, \text{depth}(u) - 2^k), u, \text{LA}(v, \text{depth}(v) - 2^k), v)$ in $O(1)$ time: we lookup the positions of $\text{LA}(u, \text{depth}(u) - 2^k) \rightsquigarrow u$ and $\text{LA}(v, \text{depth}(v) - 2^k) \rightsquigarrow v$ on the sorted list and use the range minimum query structure to calculate their longest common prefix, all in $O(1)$ time. The total space usage is $O(n)$, because every node stores one number and additionally we have a list of at most n numbers augmented with a range minimum query structure.

To construct the structure efficiently, we need to lexicographically sort all paths of length k . This can be done in $O(n)$ time for every k after observing that every path of length 2^{k+1} can be conceptually divided into two paths of length 2^k . Therefore, if we have already lexicographically sorted all paths of length 2^k , we can lexicographically sort all paths of length 2^{k+1} by sorting pairs of numbers from $[1, n]$, which are the positions of the prefix and the suffix of a longer path on the sorted list of all paths of length 2^k . With radix sorting, this takes $O(n)$ time. Then we need to compute the longest common extension of every two neighbours on the sorted list, which can be done in $O(1)$ time by using the already constructed structure for paths of length 2^k . Consequently, the total construction time is $O(n \log n)$. \square

To decrease the space usage of the structure from Lemma 2, we use the difference covers developed in Lemma 1. Intuitively, the first step is to apply the lemma with $x = \log n$ and preprocess only paths of length $2^k \log^2 n$ ending at the marked nodes. Because we have only $O(n/\log n)$ marked nodes, this requires $O(n)$ space. Then, given two paths of length ℓ , we can either immediately return their LCE using the preprocessed data, or reduce the query to computing the LCE of two paths of length at most $\log^2 n$. Using the same reasoning again with $x = \log(\log^2 n)$, we can reduce the length even further to at most $\log^2(\log^2 n)$ and so on. After $O(\log^* n)$ such reduction steps, we guarantee that the paths are of length $O(1)$, and the answer can be found naively. Formally, every step is implemented using the following lemma.

Lemma 3. *For a tree T with n nodes and a parameter b , a data structure of size $O(n)$ can be constructed in $O(n)$ time, so that given two paths of length at most b ending at $u \in T$ and $v \in T$ in $O(1)$ time we can either compute the path-path LCE or reduce the query so that the paths are of length at most $\log^2 b$.*

Proof. We apply Lemma 1 with $x = \log b$. Then, for every $k = 0, 1, \dots, \log b$ separately, we consider all paths of length $2^k \log^2 b$ ending at marked nodes. As in the proof of Lemma 2, we lexicographically sort all such paths and store the longest common extension of every two neighbours on the sorted list augmented with a range minimum query structure. Because we have only $O(n/\log b)$ marked nodes, the space decreases to $O(n)$. Furthermore, because the length of the paths is of the form $2^k \log^2 b$ (as opposed to the more natural choice of 2^k), all lists can be constructed in $O(n)$ total time by radix sorting, as a path of length $2^{k+1} \log^2 b$ ending at a marked node can be decomposed into two paths of length $2^k \log^2 b$ ending at marked nodes, because if a node is marked, so is its (x^2) -th ancestor.

Consider two paths of the same length $\ell \leq b$ ending at $u \in T$ and $v \in T$. We need to either determine their LCE, or reduce the query to determining the LCE of two paths of length at most $\log^2 b$. If $\ell \leq \log^2 b$, there is nothing to do. Otherwise, first we check if the prefixes of length $\log^2 b$ of both paths are different in $O(1)$ time. If so, we replace the paths with their prefixes of such length and we are done. Otherwise, if $\ell \leq 2 \log^2 b$ we replace the paths with their suffixes of length $\ell - \log^2 b \leq \log^2 b$ and we are done. The remaining case is that the prefixes of length $\log^2 b$ are identical and $\ell > 2 \log^2 b$. In such case, we can calculate k such that $2^k \log^2 b \leq \ell - \log^2 b < 2^{k+1} \log^2 b$. Having such k , we cover the suffixes of length $\ell - \log^2 b$ with two (potentially overlapping) paths of length exactly $2^k \log^2 b$. More formally, we create two pairs of paths:

1. $\text{LA}(u, \text{depth}(u) - 2^k \log^2 b) \rightsquigarrow u$ and $\text{LA}(v, \text{depth}(v) - 2^k \log^2 b) \rightsquigarrow v$,
2. $\text{LA}(u, \text{depth}(u) - \ell + \log^2 b) \rightsquigarrow \text{LA}(u, \text{depth}(u) - \ell + \log^2 b + 2^k \log^2 b)$ and $\text{LA}(v, \text{depth}(v) - \ell + \log^2 b) \rightsquigarrow \text{LA}(v, \text{depth}(v) - \ell + \log^2 b + 2^k \log^2 b)$.

If the paths from the first pair are different, it is enough to compute their LCE. If they are identical, it is enough to compute the LCE of the paths from the second pair. Because we can distinguish between these two cases in $O(1)$ time, we focus on computing the LCE of two paths of length $2^k \log^2 b$ ending at some u' and v' . The important additional property guaranteed by how we have defined the pairs is that the paths of length $\log^2 b$ ending at $\text{LA}(u', \text{depth}(u') - 2^k \log^2 b)$ and $\text{LA}(v', \text{depth}(v') - 2^k \log^2 b)$ are the same. Now by the properties of the difference cover we can calculate in $O(1)$ time $d \leq \log^2 b$ such that the d -th ancestors of u' and v' are marked. We conceptually slide both paths up by d , so that they both end at these marked nodes. Because of the additional property, either the paths of length $2^k \log^2 b$ ending at $\text{LA}(u', \text{depth}(u') - d)$ and $\text{LA}(v', \text{depth}(v') - d)$ are identical, or their first mismatch actually corresponds to the LCE of the original paths ending at u' and v' . These two cases can be distinguished in $O(1)$ time. Then we either use the preprocessed data to calculate the LCE in $O(1)$ time, or we are left with the suffixes of length d of the paths ending at u' and v' . But because $d \leq \log^2 b$, also in the latter case we are done. \square

We apply Lemma 3 with $b = n, \log^2 n, \log^2(\log^2 n), \dots$ terminating when $b \leq 4$. The total number of applications is just $O(\log^* n)$, because $\log^2(\log^2 z) = 4 \log^2(\log z) \leq \log z$ for z large enough⁵. Therefore, the total space usage becomes

⁵ This follows from $\lim_{z \rightarrow \infty} \frac{\log^2(\log^2 z)}{\log z} = \lim_{z \rightarrow \infty} \frac{4 \log(\log^2 z)}{\ln z} = \lim_{z \rightarrow \infty} \frac{8}{\ln z} = 0$.

$O(n \log^* n)$ and, by iteratively applying the reduction step, for any two paths of length at most n ending at given u and v we can in $O(\log^* n)$ time either compute their LCE, or reduce the query to computing the LCE of two paths of length $O(1)$, which can be computed naively in additional $O(1)$ time.

To prove Theorem 1, we need to decrease the space usage from $O(n \log^* n)$ down to $O(n)$. To this end, we create a smaller tree T' on $O(n/b)$ nodes, where $b = \log^* n$ is the parameter of the difference cover, as follows. Every marked node $u \in T$ becomes a node of T' . The parent of $u \in T$ in T' is the node corresponding in T' to the (b^2) -th ancestor of u in T , which is always marked. Additionally, we add one artificial node, which serves as the root of the whole T' , and make it the parent of all marked nodes at depth (in T) less than b^2 . Now edges of T' correspond to paths of length b^2 in T (except for the edges outgoing from the root; we will not be using them). We need to assign unique names to these paths, so that the names of two paths are equal iff the paths are the same. This can be done by traversing the suffix tree of T in $O(n)$ time. Finally, T' is preprocessed by applying Lemma 3 $O(\log^* n)$ times as described above. Because its size of T' is just $O(n/b)$, the total space usage preprocessing time is just $O(n)$ now.

To compute the LCE of two paths of length ℓ ending at $u \in T$ and $v \in T$, we first compare their prefixes of length b^2 . If they are identical, by the properties of the difference cover we can calculate $d \leq b^2$ such that the d -th ancestors of both u and v , denoted u' and v' , are marked, hence exist in T' . Consequently, if the prefixes of length $\ell - d$ of the paths are different, we can calculate their first mismatch by computing the first mismatch of the paths of length $\lfloor (\ell - d)/b^2 \rfloor$ ending at $u' \in T'$ and $v' \in T'$. This follows because every edge of T' corresponds to a path of length b^2 in T , so a path of length $\lfloor (\ell - d)/b^2 \rfloor$ in T' corresponds to a path of length belonging to $[\ell - d - b^2, \ell - d]$ in T , and we have already verified that the first mismatch is outside of the prefix of length b^2 of the original paths. Hence the first mismatch of the corresponding paths in T' allows us to narrow down where the first mismatch of the original paths in T occurs up to b^2 consecutive edges. All in all, in $O(1)$ time plus a single path-path LCE query in T' we can reduce the original query to a query concerning two paths of length at most b^2 .

The final step is to show that T can be preprocessed in $O(n)$ time and space, so that the LCE of any two paths of length at most b^2 can be calculated in $O(b)$ time. We assign unique names to all paths of length b in T , which can be again done by traversing the suffix tree of T in $O(n)$ time. More precisely, every $u \in T$ such that $\text{depth}(u) \geq b$ stores a single number, which is the name of the path of length b ending at u . To calculate the LCE of two paths of length at most b^2 ending at $u \in T$ and $v \in T$, we proceed as follows. We traverse both paths in parallel top-down moving by b edges at once. Using the preprocessed names, we can check if the first mismatch occurs on these b consecutive edges, and if so terminate. Therefore, after at most b steps we are left with two paths of length at most b , such that computing their LCE allows us to answer the original query. But this can be calculated by naively traversing both paths in parallel top-down. The total query time is $O(b)$.

To summarize, the total space and preprocessing time is $O(n)$ and the query time remains $O(\log^* n)$, which proves Theorem 1.

5 Path-Tree LCE

In this section we prove the following theorem.

Theorem 2. *For a tree T with n nodes, a data structure of size $O(n)$ can be constructed in $O(n)$ time to answer path-tree LCE queries in $O((\log \log n)^2)$ time.*

The idea is to apply the difference covers recursively with the following lemma.

Lemma 4. *For a tree T with n nodes and a parameter b , a data structure of size $O(n)$ can be constructed in $O(n \log n)$ time, so that given a path of length $\ell \leq b$ ending at $u \in T$ and a subtree rooted at $v \in T$ we can reduce the query in $O(\log \log n)$ time so that the path is of length at most $b^{4/5}$.*

Proof. The first part of the structure is designed so that we can detect in $O(1)$ time if the path-tree LCE is of length at most $b^{4/5}$. We consider all paths of length exactly $b^{4/5}$ in the tree. We assign names to every such path, so that testing if two paths are identical can be done by looking at their names. Then, for every node w we gather all paths of length $b^{4/5}$ starting at w (i.e., $w \rightsquigarrow v$, where $w = \text{LA}(v, \text{depth}(v) - b^{4/5})$) and store their names in a perfect hash table, where every name is linked to the corresponding node w . This allows us to check if the answer is at least $b^{4/5}$ by first looking up the name of the prefix of length $b^{4/5}$ of the path, and then querying the perfect hash table kept at v . If the name does not occur there, the answer is less than $b^{4/5}$ and we are done. Otherwise, we can move by $b^{4/5}$ down, i.e., decrease ℓ by $b^{4/5}$ and replace v with its descendant of distance $b^{4/5}$.

The second part of the structure is designed to work with the marked nodes. We apply Lemma 1 with $x = b^{2/5}$ and consider *canonical paths* of length $i \cdot x^2$ in the tree, where $i = 1, 2, \dots, \sqrt{x}$, ending at marked nodes. The total number of such paths is $O(n/\sqrt{x})$, because every marked node is the endpoint of at most \sqrt{x} of them. We lexicographically sort all canonical paths and store the longest common extension of every two neighbours on the global sorted list augmented with a range minimum query structure. Also, for every marked node v and every $i = 1, 2, \dots, x$, we save the position of the path $\text{LA}(v, \text{depth}(v) - i \cdot x^2) \rightsquigarrow v$ on the global sorted list. Additionally, at every node u we gather all canonical paths starting there, i.e., $u \rightsquigarrow v$ such that $\text{LA}(v, \text{depth}(v) - i \cdot x^2) = u$ for some $i = 1, 2, \dots, x$, sort them lexicographically and store on the local sorted list of u . Every such path is represented by a pair (u, i) . The local sorted list is augmented with a predecessor structure storing the positions on the global sorted list.

Because we have previously decreased ℓ and replaced v , now by the properties of the difference cover we can find $d \leq x^2$ such that the $(\ell+d)$ -th ancestor of u and the d -th ancestor of v are marked, and then increase ℓ by d and replace v by its d -th ancestor. Consequently, from now on we assume that both $\text{LA}(u, \text{depth}(u) - \ell)$ and v are marked.

Now we can use the second part of the structure. If $\ell \leq b^{4/5}$, there is nothing to do. Otherwise, the prefix of length $\lfloor \ell/x^2 \rfloor \cdot x^2$ of the path is a canonical path (because $\ell \leq \sqrt{x} \cdot x^2$), so we know its position on the global sorted list. We query the predecessor structure stored at v with that position to get the lexicographical

predecessor and successor of the prefix among all canonical paths starting at v . This allows us to calculate the longest common extension p of the prefix and all canonical paths starting at v by taking the maximum of the longest common extension of the prefix and its predecessor, and the prefix and its successor. Now, because canonical paths are all paths of the form $i \cdot x^2$, the length of the path-tree LCE cannot exceed $p + x^2$. Furthermore, with a level ancestor query we can find v' such that the paths $\text{LA}(u, \text{depth}(u) - \ell) \rightsquigarrow \text{LA}(u, \text{depth}(u) - \ell + p)$ and $v \rightsquigarrow v'$ are identical. Then, to answer the original query, it is enough to calculate the path-tree LCE for $\text{LA}(u, \text{depth}(u) - \ell + p) \rightsquigarrow \text{LA}(u, \text{depth}(u) - \ell + \min(\ell, p + x^2))$ and the subtree rooted at v' . Therefore, in $O(\log \log n)$ time we can reduce the query so that the path is of length at most $x^2 = b^{4/5}$ as claimed.

To achieve $O(n)$ construction time, we need to assign names to all paths of length $b^{4/5}$ in the tree, which can be done in $O(n)$ by traversing the suffix tree of T . We would also like to lexicographically sort all canonical paths, but this seems difficult to achieve in $O(n)$. Therefore, we change the lexicographical order as follows: we assign names to all canonical paths of length exactly x^2 , so that different paths get different names and identical paths get identical names (again, this can be done in $O(n)$ time by traversing the suffix tree). Then we treat every canonical path of length $i \cdot x^2$ as a sequence of consisting of i names, and sort these sequences lexicographically in $O(n)$ time with radix sort. Even though this is not the lexicographical order, the canonical paths are only used to approximate the answer up to an additive error of x^2 , and hence such modification is still correct. \square

We apply Lemma 4 with $b = n, n^{4/5}, n^{(4/5)^2}, \dots, 1$. The total number of applications is $O(\log \log n)$. Therefore, the total space usage becomes $O(n \log \log n)$, and by applying the reduction step iteratively, for any path of length n ending at u and a subtree rooted at v we can compute the path-tree LCE in $O((\log \log n)^2)$ time. The total construction time is $O(n \log \log n)$.

To prove Theorem 2, we need to decrease the space usage and the construction time. The idea is similar to the one from Section 4: we create a smaller tree T' on $O(n/b)$ nodes, where $b = \log \log n$ is the parameter of the difference cover. The edges of T' correspond to paths of length b^2 in T . We preprocess T' as described above, but because its size is now just $O(n/b)$, the preprocessing time and space become $O(n)$.

To compute the path-tree LCE for a given path of length ℓ ending at u and a subtree rooted at v , we first check if the answer is at least b^2 . This can be done in $O(\log \log n)$ time by preprocessing all paths of length b^2 in T , as done inside Lemma 4 for paths of length $b^{4/5}$. If so, we can decrease ℓ and replace v with its descendant, so that both $\text{LA}(u, \text{depth}(u) - \ell)$ and v are marked, hence exist in T' . Then we use the structure constructed for T' to reduce the query, so that the path is of length at most b^2 . Therefore, it is enough how to answer a query, where a path is of length at most b^2 , in $O(\log \log n)$ time after $O(n)$ time and space preprocessing.

The final step is to preprocess T in $O(n)$ time and space, so that the path-tree LCE of a path of length at most b^2 and any subtree can be computed in $O(b)$ time. We assign unique names to all paths of length b in T . Then, for every u we gather the names of all paths $u \rightsquigarrow v$ of length b in a perfect hash table. To

calculate the path-tree LCE, we traverse the path top-down while tracing the corresponding node in the subtree. Initially, we move by b edges by using the perfect hash tables. This allows us to proceed as long as the remaining part of the LCE is at least b . Then, we traverse the remaining part consisting of at most b edges naively. In total, this takes $O(b)$ time. The space is clearly $O(n)$ and the preprocessing requires constructing the perfect hash tables, which can be done in $O(n)$ time.

5.1 Lower Bound

In this section, we prove that any path-tree LCE structure of size $O(n \text{ polylog}(n))$ must necessarily use $\Omega(\log \log n)$ time to answer queries. As shown by Pătraşcu and Thorup [22], for $U = n^2$ any predecessor structure consisting of $O(n \text{ polylog}(n))$ words needs $\Omega(\log \log n)$ time to answer queries, assuming that the word size is $\Theta(\log n)$. In the full version of this paper, we show the following reduction, which implies the aforementioned lower bound.

Theorem 3. *For any $\epsilon > 0$, given an LCE_{PT} structure that uses $s(n) = \Omega(n)$ space and answers queries in $q(n) = \Omega(1)$ time we can build a predecessor structure using $O(s(2U^\epsilon + n \log |U|))$ space and $O(q(2U^\epsilon + n \log |U|))$ query time for any $S \subseteq [0, U)$ of size n .*

By applying the reduction with $U = n^2$ and $\epsilon = 1/2$, we get that an LCE_{PT} structure using $O(n \text{ polylog}(n))$ space and answering queries in $o(\log \log n)$ time implies a predecessor structure using $O(n \text{ polylog}(n))$ space and answering queries in $o(\log \log(n))$ time, which is not possible.

6 Tree-Tree LCE

The *set intersection problem* is defined as follows. Given a family $\mathcal{S} = \{S_1, \dots, S_k\}$ of sets of total size $n = \sum_{i=1}^k |S_i|$ the goal is to preprocess \mathcal{S} to answer queries: given two sets S_i and S_j determine if $S_i \cap S_j = \emptyset$. The set intersection problem is widely believed to require superlinear space in order to support fast queries. A folklore conjecture states that for sets of size polylogarithmic in k , supporting queries in constant time requires $\tilde{\Omega}(k^2)$ space [21] (see also [8]).

We now consider the LCE_{TT} problem. In the full version of this paper we show that the problem is set intersection hard and give a time-space trade-off as stated by the following theorems:

Theorem 4. *Let T be a tree with n nodes. Given an LCE_{TT} data structure that uses $s(n)$ space and answers queries in $q(n)$ time we can build a set intersection data structure using $O(s(n))$ space and $O(q(n))$ query time, for input sets containing $O(n)$ elements.*

Theorem 5. *For a tree T with n nodes and a parameter τ , $1 \leq \tau \leq n$, a data structure of size $O(n\tau)$ can be constructed in $O(n\tau)$ time to answer tree-tree LCE queries in $O(n/\tau)$ time.*

References

1. S. Alstrup and J. Holm. Improved algorithms for finding level ancestors in dynamic trees. In *Proc. 27th ICALP*, pages 73–84, 2000.
2. A. Amir, M. Lewenstein, and E. Porat. Faster algorithms for string matching with k mismatches. *J. Algorithms*, 50(2):257–275, 2004.
3. H. Bannai, P. Gawrychowski, S. Inenaga, and M. Takeda. Converting SLP to LZ78 in almost linear time. In *Proc. 24th CPM*, pages 38–49, 2013.
4. M. A. Bender and M. Farach-Colton. The LCA problem revisited. In *Proc. of 4th LATIN*, pages 88–94, 2000.
5. M. A. Bender and M. Farach-Colton. The level ancestor problem simplified. *Theoret. Comput. Sci.*, 321(1):5 – 12, 2004.
6. O. Berkman and U. Vishkin. Finding level-ancestors in trees. *J. Comput. System Sci.*, 48(2):214 – 230, 1994.
7. D. Breslauer. The suffix tree of a tree and minimizing sequential transducers. *Theoret. Comput. Sci.*, 191(1–2):131 – 144, 1998.
8. H. Cohen and E. Porat. Fast set intersection and two-patterns matching. *Theor. Comput. Sci.*, 411(40–42):3795–3800, 2010.
9. R. Cole and R. Hariharan. Approximate string matching: A simpler faster algorithm. *SIAM J. Comput.*, 31(6):1761–1782, 2002.
10. P. F. Dietz. Finding level-ancestors in dynamic trees. In *Proc. 2nd WADS*, pages 32–40, 1991.
11. J. Fischer and V. Heun. Space-efficient preprocessing schemes for range minimum queries on static arrays. *SIAM J. Comput.*, 40(2):465–492, 2011.
12. M. L. Fredman, J. Komlos, and E. Szemerédi. Storing a sparse table with $O(1)$ worst case access time. In *Proc. 23rd FOCS*, pages 165–169, Nov 1982.
13. R. F. Geary, R. Raman, and V. Raman. Succinct ordinal trees with level-ancestor queries. *ACM Trans. Algorithms*, 2(4):510–534, 2006.
14. D. Gusfield. *Algorithms on strings, trees, and sequences: computer science and computational biology*. Cambridge, 1997.
15. D. Gusfield and J. Stoye. Linear time algorithms for finding and representing all the tandem repeats in a string. *J. of Comput. Syst. Sci.*, 69(4):525 – 546, 2004.
16. D. Harel and R. E. Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM J. Comput.*, 13(2):338–355, 1984.
17. S. R. Kosaraju. Efficient tree pattern matching. In *Proc. of 30th FOCS*, pages 178–183, 1989.
18. G. M. Landau, E. W. Myers, and J. P. Schmidt. Incremental string comparison. *SIAM J. Comput.*, 27(2):557–582, 1998.
19. G. M. Landau and U. Vishkin. Fast parallel and serial approximate string matching. *J. Algorithms*, 10:157–169, 1989.
20. M. G. Main and R. J. Lorentz. An $O(n \log n)$ algorithm for finding all repetitions in a string. *J. Algorithms*, 5(3):422 – 432, 1984.
21. M. Pătrașcu and L. Roditty. Distance oracles beyond the Thorup–Zwick bound. In *Proc. 51st IEEE FOCS*, pages 815–823, 2010. To appear.
22. M. Pătrașcu and M. Thorup. Time-space trade-offs for predecessor search. In *Proc. of 38th STOC*, pages 232–240, 2006.
23. M. Ružić. Algorithms for deterministic construction of efficient dictionaries. In *Proc. of 12th ESA*, pages 592–603, 2004.
24. T. Shibuya. Constructing the suffix tree of a tree with a large alphabet. In *Proc. 10th ISAAC*, pages 225–236, 1999.
25. P. van Emde Boas, R. Kaas, and E. Zijlstra. Design and implementation of an efficient priority queue. *Mathematical Systems Theory*, 10:99–127, 1977.