

Optimal Distance Labeling Schemes for Trees

Ofer Freedman^{*1}, Paweł Gawrychowski^{*1}, Patrick K. Nicholson², and Oren Weimann^{*1}

¹University of Haifa, Israel

²Bell Labs, Dublin, Ireland

Abstract

Labeling schemes seek to assign a short label to each node in a network, so that a function on two nodes (such as distance or adjacency) can be computed by examining their labels alone. For the particular case of trees, following a long line of research, optimal bounds (up to low order terms) were recently obtained for adjacency labeling [FOCS '15], nearest common ancestor labeling [SODA '14], and ancestry labeling [SICOMP '06]. In this paper we obtain optimal bounds for distance labeling. We present labels of size $1/4 \log^2 n + o(\log^2 n)$, matching (up to low order terms) the recent $1/4 \log^2 n - O(\log n)$ lower bound [ICALP '16].

Prior to our work, all distance labeling schemes for trees could be reinterpreted as *universal trees*. A tree T is said to be universal if any tree on n nodes can be found as a subtree of T . A universal tree with $|T|$ nodes implies a distance labeling scheme with label size $\log |T|$. In 1981, Chung et al. proved that any distance labeling scheme based on universal trees requires labels of size $1/2 \log^2 n - \log n \cdot \log \log n + O(\log n)$. Our scheme is the first to break this lower bound, showing a separation between distance labeling and universal trees.

The $\Theta(\log^2 n)$ barrier for distance labeling in trees has led researchers to consider distances bounded by k . The size of such labels was improved from $\log n + O(k\sqrt{\log n})$ [WADS '01] to $\log n + O(k^2 \log(k \log n))$ [SODA '03] and finally to $\log n + O(k \log(k \log(n/k)))$ [PODC '07]. We show how to construct labels whose size is the minimum between $\log n + O(k \log((\log n)/k))$ and $O(\log n \cdot \log(k/\log n))$. We complement this with almost tight lower bounds of $\log n + \Omega(k \log(\log n/(k \log k)))$ and $\Omega(\log n \cdot \log(k/\log n))$. Finally, we consider $(1 + \varepsilon)$ -approximate distances. We show that the recent labeling scheme of [ICALP '16] can be easily modified to obtain an $O(\log(1/\varepsilon) \cdot \log n)$ upper bound and we prove a matching $\Omega(\log(1/\varepsilon) \cdot \log n)$ lower bound.

^{*}The research was supported in part by Israel Science Foundation grant 794/13.

1 Introduction

Labeling schemes seek to assign a short label to each vertex in a network, so that a function on two nodes (such as distance or adjacency) can be computed by examining their labels alone. This is particularly desirable in distributed settings, where nodes are often processed using only some locally stored data. Recently, with the rise in popularity of distributed computing platforms such as Spark and Hadoop, labeling schemes have found renewed interest. Indeed, the goal of minimizing the size of the maximal label has been the subject of a great deal of recent research [1, 5–10, 17, 18, 27]. For the particular case of trees, the functions that have been studied are distance [5, 8, 20, 21, 26], adjacency [6, 11, 13], nearest common ancestor [9, 17], and ancestry [1, 18] (a recent survey of these results can be found here [29]). Tree labeling schemes have recently found new uses in large scale graph processing. For example, distance oracles for general graphs use distance labelings for spanning trees rooted at judiciously chosen vertices [2–4].

Universal trees. A particularly clean way of looking at labeling schemes is through *universal graphs*. A graph G is said to be *universal* for a given family of graphs, if every graph in the family is an induced subgraph of G . Similarly, a tree T is said to be universal for *all* trees on n nodes if any tree on n nodes can be found as a subtree of T . For *adjacency labeling in graphs*, Kannan et al. [24] observed that if a family of graphs has a universal graph with $|G|$ vertices then it has an adjacency labeling scheme with label size $\log |G|$, and vice versa. For *distance labeling in trees*, until the present work, this statement was only known to be true in one direction. Namely, a universal tree T of all trees on n nodes implies a distance labeling scheme with label size $\log |T|$. We prove that the converse is in fact not true.

The use of universal trees is powerful, but it is limited. Already 50 years ago, Goldberg and Livshits [22] showed how to construct a universal tree T that is of size $|T| = n^{(\log n - 2 \log \log n + O(1))/2}$ which was shown by Chung et al. [14] to be the smallest possible up to the $O(1)$ error term. This shows the first limitation of using universal trees for distance labeling: there is a lower bound of $\log |T| = 1/2 \log^2 n - \log n \cdot \log \log n + O(\log n)$ on the label size. The second limitation is the query time. The universal tree construction of Goldberg and Livshits was given before labeling schemes were ever invented. Of course, one could naively use their universal tree T for distance labeling of an arbitrary tree on n nodes by finding its isomorphic subtree in T and assigning labels which are just the IDs of the nodes in T . However, such a non-algorithmic labeling would require prohibitive query time and space since T needs to be computed. This latter limitation was overcome by algorithmic labeling schemes achieving logarithmic query time: An upper bound of $O(\log^2 n)$ bits on the label size was first shown by Peleg [26] and a lower bound of $1/8 \log^2 n - O(\log n)$ bits was shown by Gavaille et al. [21]. Very recently, Alstrup et al. [8] improved the lower bound to $1/4 \log^2 n - O(\log n)$ and observed that the upper bound can be improved to $1/2 \log^2 n + O(\log n)$ with a somewhat straightforward use of a nearest common ancestor labeling scheme.

All the above labeling schemes can be reinterpreted as building a universal tree, and are therefore subject to the $1/2 \log^2 n - \log n \cdot \log \log n + O(\log n)$ lower bound of Chung et al. In other words, the scheme of Alstrup et al. is optimal (up to low order terms) amongst all schemes that translate to universal trees. To see why the scheme of Alstrup et al. indeed translates to a universal tree, we show in Section 3.6 that their scheme can be casted as a level-ancestor scheme and we show in Section 3.5 that every level-ancestor scheme translates to a universal tree.

We give the first distance labeling scheme that does not translate to a universal tree. This enables us to circumvent the Chung et al. [14] lower bound for labels based on universal trees and to match the general lower bound of Alstrup et al. [8]. Namely, in Section 3 we prove the following:

Theorem 1.1. *There is a scheme for tree distance labeling with $1/4 \log^2 n + o(\log^2 n)$ bit labels and constant query time.*

The above theorem means that universal trees capture more than is required for distance labeling. To illustrate this, we need to describe the related problem of level-ancestor labelings.

Labeling schemes for level-ancestor. In this problem, we are given a rooted tree and seek to assign labels so that we can compute (the label of) any k -th ancestor of a node from its label alone. Notice that here a query receives a single label and a value k , and that all labels must be distinct (no scheme which uses the same label twice can be correct).

It is not hard to see that labels supporting level-ancestor queries can be used to answer distance queries. Thus, any lower bound for tree distance labeling immediately applies to level-ancestor labeling, but the converse is not true. Nevertheless, it turns out that all previous distance labeling schemes are also level-ancestor schemes. Like the labeling scheme of Alstrup et al. [8], our scheme is also based on a *heavy path decomposition* of the tree, which can be seen as a way of transforming an arbitrary tree into an edge-weighted tree of logarithmic depth. However, while the labels in [8] store the weights of every edge on the path to the root (thus allowing for level-ancestor queries), we show that it is possible to carefully distribute the bits between the labels so that the distance can be computed given any pair of labels, yet a single label is not enough to extract the level-ancestors.

We determine this separation between tree distance labeling and level-ancestor labeling by proving that labeling for distances is roughly half as expensive as labeling for level-ancestors:

Theorem 1.2. *Any scheme for level-ancestor labeling must use at least $1/2 \log^2 n - \log n \log \log n$ bits for the maximum length label.*

We prove the above theorem in Section 3.5 by showing that, as opposed to distance labeling, no level-ancestor labeling scheme can do better than the one based on universal trees. Namely, we prove that any level-ancestor labeling scheme with labels of length L implies a universal rooted tree of size $O(2^L)$, and then invoke the known lower bound for universal trees [14, 22]. In particular, it means that for level-ancestor queries, the scheme of Alstrup et al. [8] is optimal (after some modifications described in Section 3.6).

Labeling schemes for bounded distances. The $\Theta(\log^2 n)$ barrier on distance labeling in trees has initiated a line of research that improves the label size when the distances are bounded: In *k -distance labeling*, we are given the labels of u and v and need to decide if the length of the u -to- v path is at most k , and if so return it. For $k = 1$, this is exactly adjacency labeling, which was recently shown by Alstrup et al. [6] to require only $\log n + O(1)$ bits. For $k \geq 2$, this was first considered by Kaplan and Milo [25] who showed how to construct labels of length $\log n + O(k\sqrt{\log n})$. The query time was not explicitly specified in their implementation, but appears to be $O(k)$. A shorter label of $\log n + O(k^2 \log(k \log n))$ bits was then given by Alstrup, Bille, and Rauhe [5] who also proved that any scheme for $k \geq 2$ (i.e., any scheme that is able to answer “ancestor or sibling” queries) requires $\log n + \Omega(\log \log n)$ bits. Hence the $O(\log \log n)$ addend cannot be avoided, but it remained unclear what should be the exact dependency on k nor the query time (Alstrup, Bille, and Rauhe considered constant k in which case their bounds are tight and their $O(k^2)$ query time is constant). The labeling scheme of Alstrup, Bille, and Rauhe was then improved by Gavaille and Labourel [20] who presented a bound of $\log n + O(k \log(k \log(n/k)))$ bits and $O(k)$ query time.

In Section 4 we show how to construct a labeling scheme with improved label size and constant query time, and prove an almost matching lower bound. Formally, we prove:

Theorem 1.3. *For $k < \log n$, there is a k -distance labeling scheme with labels of length $\log n + O(k \log((\log n)/k))$ bits, and any such scheme requires $\log n + \Omega(k \log(\log n/(k \log k)))$ bits.*

For $k \geq \log n$, there is a k -distance labeling scheme with labels of length $O(\log n \cdot \log(k/\log n))$ bits, and any such scheme requires $\Omega(\log n \cdot \log(k/\log n))$ bits. In both cases, the query time is constant.

For the upper bound, our starting point is the scheme of Alstrup, Bille, and Rauhe [5]. We observe that, instead of storing the same information for each of the nearest k heavy paths above a node, it is possible to store all information for the topmost of these heavy paths and less information for all the rest. To improve the query time, we show that only a subtle change is needed in the definition of the so-called significant preorder numbers. The new definition retains all the nice properties of the previous while being much easier to operate on. Our constant query time assumes the standard word-RAM model with word size $\Omega(\log n)$.

For the lower bounds we take two different approaches. For $k < \log n$, we show how to construct a family of trees such that, in any k -distance labeling scheme, different trees can share some labels but every tree has to introduce many additional unique labels. For $k \geq \log n$, we use the clever lower bound technique from (unbounded) distance labelings, that was introduced by Gavoille et al. [21] and refined by Alstrup et al. [8]. It is based on constructing a *weighted* almost complete binary tree, where all the leaves are at the same distance from the root. After arguing that the labels of nodes in such a tree must be long, the weights are removed by subdividing edges while not increasing the size of the tree by too much. We show that only a small tweak is required to this known lower bound for distance labeling in order to get a lower bound for k -distance labeling.

Labeling schemes for approximate distances. Finally, we consider $(1 + \varepsilon)$ -approximate distance labeling, where given the labels of u and v we need to output a value that is at least $d(u, v)$ and at most $(1 + \varepsilon) \cdot d(u, v)$. For the case $\varepsilon = 1/\log n$, Gavoille et al. [19] proved a tight bound of $\Theta(\log(1/\varepsilon) \cdot \log n)$. Very recently, Alstrup et al. [8] considered the general trade-off and designed, for any constant $\varepsilon \leq 1$, an $O(\log n)$ bit labeling scheme. In Section 5 we show that their solution can be easily made to produce labels of size $O(\log(1/\varepsilon) \cdot \log n)$ and that this is the best possible:

Theorem 1.4. *For any $\varepsilon \leq 1$, there is a $(1 + \varepsilon)$ -approximate distance labeling scheme with labels of length $O(\log(1/\varepsilon) \cdot \log n)$ bits, and any such scheme requires $\Omega(\log(1/\varepsilon) \cdot \log n)$ bits.*

The lower bound is obtained by reducing exact distance labeling to $(1 + \varepsilon)$ -approximate distance labeling. This is achieved by appropriately stretching the lengths of the edges in the lower bound instances of Gavoille et al. [21]. For the upper bound, we slightly modify the scheme of Alstrup et al. [8], which originally stored a sequence of integers using simple unary encoding. Such an encoding requires $O(1/\varepsilon \cdot \log n)$ bits. We show that with a more complicated binary encoding we can obtain a scheme with $O(\log(1/\varepsilon) \cdot \log n)$ bits and a constant query time. We conclude this section with the following table summarizing our contribution.

Label type	Upper bound	Lower bound	
Exact	$1/4 \log^2 n + o(\log^2 n)$	$1/4 \log^2 n - O(\log n)$ [8]	
Approximate	$O(\log(1/\varepsilon) \cdot \log n)$	$\Omega(\log(1/\varepsilon) \cdot \log n)$	
k -distance	$k \geq \log n$	$O(\log n \cdot \log \frac{k}{\log n})$	$\Omega(\log n \cdot \log \frac{k}{\log n})$
	$k < \log n$	$\log n + O(k \log \frac{\log n}{k})$	$\log n + \Omega(k \log \frac{\log n}{k \log k})^1$

¹ This lower bound only holds for $k = o(\frac{\log n}{\log \log n})$.

2 Preliminaries

We consider a rooted tree T , or we arbitrarily root it. We denote the root by $\text{root}(T)$, and the distance between node v to $\text{root}(T)$ by $\text{root-distance}(v)$. We denote the subtree rooted at u as T_u , and the number of nodes of T by $|T|$, or simply n if T is known from the context. For two nodes u, v , we denote their distance by $d(u, v)$, their nearest common ancestor by $\text{NCA}(u, v)$. First, observe that:

$$d(u, v) = \text{root-distance}(u) + \text{root-distance}(v) - 2 \cdot \text{root-distance}(\text{NCA}(u, v)).$$

This means that a labeling scheme for $\text{root-distance}(\text{NCA}(u, v))$ that assumes $u \neq v$ can actually be used for $d(u, v)$ queries with only $O(\log n)$ additional bits to the label size (the additional bits are simply the distance to the root). Next, observe that although our input tree is unweighted (i.e., all edges have weight 1), if our distance labeling scheme can handle edges-weights in $\{0, 1\}$ then we can assume the input tree is binary and that the queries are on leaves only. This can be achieved by connecting every internal node u to a leaf node u_ℓ with an edge of weight 0, and then standardly binarizing the tree (by inserting $O(n)$ intermediate nodes with edge-weights 0 connecting them).

Heavy path decompositions. We apply a variant of heavy path decompositions [31]. We start at the root of the tree T and repeatedly descend from the current node to its (unique) child whose subtree is of size at least $|T|/2$ as long as possible, that is, we terminate when there is no such child. Note that this is different than the more common versions in which we descend from the current node u to its child v with the largest subtree until (depending on the version) we reach a leaf or $|T_v| < |T_u|/2$. This gives us a *heavy path* P starting at $\text{root}(T)$ and many subtrees hanging off the heavy path. We call the edges of P *heavy*, and all other edges outgoing from the nodes of P *light*. The construction is then applied recursively to all subtrees hanging off the heavy path. In the end, each node $u \in T$ has at most one heavy child, denoted $\text{heavy}(u)$, so we obtain a decomposition of T into disjoint heavy paths (some of which consist of a single node). The light depth of a node $u \in T$, denoted $\text{lightdepth}(u)$, is the number of light edges on the path from u to the root and is at most $\log n$ [30]. We order the children of every node u so that $\text{heavy}(u)$ is the rightmost child and assign preorder numbers $\text{pre}(u)$ to every node u . Then, for any node $v \in T_u$, we have that $\text{pre}(v) \in [\text{pre}(u), \text{pre}(u) + |T_u|]$.




Figure 1: On the left, a heavy path decomposition of a binary tree T . Light nodes are white, and heavy nodes are black. The heavy edges are solid, the light edges are dashed, and the exceptional edges are dashed and hollowed. On the right, the collapsed tree $\mathcal{C}(T)$.

The collapsed tree. Given the heavy path decomposition of a *binary* tree T , we define its *collapsed tree*, denoted $\mathcal{C}(T)$, whose nodes correspond to heavy paths in T . The heavy path

starting at $\text{root}(T)$ corresponds to the root of the collapsed tree. Every light edge hanging off this heavy path corresponds to an edge outgoing from the root of $\mathcal{C}(T)$, and so on. The children of every node in $\mathcal{C}(T)$ are ordered according to the top-to-bottom order on the hanging subtrees (i.e, if two subtrees connect to the same heavy path P then the one connecting at a lower depth is to the left of the other). Since T is binary, ties can only happen at the last node of the heavy path P , in which case we set the right subtree to be the subtree of maximum size, and call the light edge branching to the right subtree the *exceptional* edge associated with heavy path P . See Figure 1 (right). Note that the height of the collapsed tree is at most $\log n$.

Every heavy path P in T is associated with a node u' in $\mathcal{C}(T)$ and every node $u \in P$ is said to be *associated* with u' . We refer to the node $u \in P$ closest to the root of T as the *head* of P and denote it as $\text{head}(P)$ or $\text{head}(u')$. We use $\text{lightdepth}(u, v)$ to denote $\text{lightdepth}(\text{NCA}(u, v))$. Finally, we say that u *dominates* v if the inorder number of u 's associated node in $\mathcal{C}(T)$ is smaller than that of v 's. Observe that (1) If the $\text{NCA}(u, v)$ -to- u path in T starts with a light edge and the $\text{NCA}(u, v)$ -to- v path starts with a heavy edge then u dominates v , and (2) If both these paths start with a light edge then the dominated vertex is the one whose path starts with the exceptional edge.

Labeling schemes for NCA. A nearest common ancestor scheme assigns a unique label to every node, so that given the labels of nodes u, v we can return the label of $\text{NCA}(u, v)$. Alstrup et al. [9] design such a scheme with labels of length $O(\log n)$ bits. They use a heavy path decomposition that slightly differs from ours, but it can be easily verified that the following lemma still holds:

Lemma 2.1 ([8, 9]). *There is an NCA labeling scheme with label size $O(\log n)$ bits, which given the labels of u and v returns the label of $\text{NCA}(u, v)$ as well as $\text{lightdepth}(u, v)$ in constant time.*

Encoding integers. To store a single integer x , we use Elias δ codes [16] that require $\log x + O(\log \log x)$ bits. This encoding is *self-delimiting*, meaning we can concatenate multiple variable-length values into a single label in a way that each individual value can be decoded later. To store a monotone sequence of integers we use the following lemma:

Lemma 2.2. *A monotone sequence of s integers in $[0, M]$ can be encoded with $O(s \cdot \max\{1, \log \frac{M}{s}\})$ bits, so that we can:*

- (1) *extract the k^{th} number in the sequence,*
- (2) *find the position of the successor of a given integer in the sequence,*
- (3) *given the representation of two sequences, find the longest common suffix of two specified prefixes.*

The first operation takes constant time, and the second and third take constant time if both s and M are $O(\log n)$.

Proof. Let the sequence be $0 \leq x_1 \leq x_2 \leq \dots \leq x_s \leq M$. The encoding consists of x_1 and the differences $x_2 - x_1, x_3 - x_2, \dots, x_s - x_{s-1}$. Each number is encoded using the Elias γ code, so the total size of the encoding becomes $O(s + \sum_{i=1}^s \log(x_i - x_{i-1}))$, where $x_0 = 0$. By Jensen's inequality, this is maximized when all numbers are equal, so the total size of the encoding is $L = O(s \cdot \max\{1, \log \frac{M}{s}\})$.

To provide constant time access to every x_i , we need to store some auxiliary data. We partition the universe $[0, M]$ into blocks of length $b = \frac{M}{s}$. For each x_i , we store $x_i \bmod b$. This is done by reserving $\lceil \log b \rceil$ bits for every $i = 1, 2, \dots, s$ and arranging them one after another. We also store b encoded using the Elias γ code, so that in constant time we can calculate where the $\lceil \log b \rceil$ bits storing $x_i \bmod b$ are. This takes $O(\log b + s + s \log b) = O(L)$ space so far. It

remains to show how to encode $y_i = x_i \text{ div } b$. Notice that $0 \leq y_1 \leq y_2 \leq \dots \leq y_s \leq s$, so this is a monotone sequence of s integers from $[0, s]$. We encode it with a single bit vector of length at most $2s$, which is the concatenation of $0^{y_i - y_{i-1}}1$ for $i = 1, 2, \dots, s$ (and $y_0 = 0$). Then, to extract y_i we need to find the position p of the i^{th} bit set to 1 in the bit vector and then return $p - i + 1$. By augmenting the bit vector with a select structure of Clark [15, Chapter 2.2], which takes $o(s)$ additional bits of space, we can retrieve the i^{th} bit set to 1 in constant time. Thus, in $O(L)$ additional space we can encode $x_i \bmod b$ and $x_i \text{ div } b$, and then recover x_i in constant time.

To provide constant time successor queries (when both s and M are $O(\log n)$), we remove all duplicates and store the resulting sequence $y_1 < y_2 < \dots < y_r$ in an additional predecessor structure from the second branch of Pătraşcu and Thorup [28]. This structure uses $O(r \cdot \log M)$ bits and answers queries in $O(\log \frac{\log M}{\log \log n}) = O(1)$ time. The space can be actually improved to $O(r \cdot \log \frac{M}{r}) = O(s \cdot \max\{1, \log \frac{M}{s}\})$ as explained in detail by Belazzougui and Navarro [12].

Finally, to compute the longest common suffixes of two specified prefixes given the encodings of $x_1 \leq x_2 \leq \dots \leq x_s$ and $y_1 \leq y_2 \leq \dots \leq y_s$, we observe that for $s, M = O(\log n)$ the encodings fit in a constant number of machine words. Hence, we can first shift both encodings (in constant time) to reduce the problem to computing the longest common suffix. First, we check if $x_s = y_s$. If not, we are done. Otherwise, we only need to find the longest common suffix of the sequences of differences. This can be done by first calculating the longest common suffix of their encodings, and then counting how many differences have their encodings fully in the common suffix. The former can be done in constant time using the standard word-RAM operations. The latter can be done by storing an additional bit vector of length L , where we mark the starting position of the encoding of each x_i with a bit set to 1. The bit vector is augmented with the rank structure of Jacobson [23], which takes $o(L)$ additional bits and allows us to count bits set to 1 in any prefix in constant time. \square

(h, M) -trees. To obtain a lower bound for distance labeling, Gavaille et al. [21] consider a family of rooted binary trees called (h, M) -trees. The trees are weighted and the weight of every edge is in $[0, M]$. For $h = 0$ the tree is a single node. For $h \geq 1$, the tree consists of a root connected to its single child with an edge of length $M - x$ for some $x \in [0, M)$, and the child is connected to two (possibly different) $(h - 1, M)$ -trees with edges of length x . See Figure 2. A lower bound for tree distance labeling is implied by the following lemma:




Figure 2: A $(3, M)$ -tree, where $x_1, \dots, x_7 \in [0, M)$.

Lemma 2.3 ([21]). *For $h \geq 1$ and $M \geq 2$, any scheme for distance labeling in (h, M) -trees requires labels of at least $h/2 \cdot \log M$ bits, even if we only query leaves.*

3 Distance Labeling

In this section we prove Theorem 1.1. In Section 3.1 we review the labeling scheme framework of the existing solutions (in a slightly different way), and in Section 3.2 we describe our improved solution and its analysis.

3.1 Distance Arrays

We now review the general framework for distance labeling. For each node $u \in T$, consider the set of light edges $\ell_1(u), \dots, \ell_k(u)$ along the root-to- u path. For any light edge e in the collapsed tree $\mathcal{C}(T)$ branching from u' to its child v' let $\mathbf{d}(e) = \mathbf{d}(\text{head}(u'), \text{head}(v'))$. That is, the distance along the heavy path represented by u' to the endpoint where the light edge branches and to its other end. Let $D(u)$ denote the list $[\mathbf{d}(\ell_1(u)), \dots, \mathbf{d}(\ell_k(u))]$, which we call the *distance array* of u . The next lemma shows that designing an efficient distance labeling scheme boils down to efficiently encoding distance arrays.

Lemma 3.1. *If we can access the elements of the distance arrays $D(u)$ and $D(v)$ then with additional $O(\log n)$ bits we can compute $\mathbf{d}(u, v)$.*

Proof. We first describe the additional $O(\log n)$ bits. They are composed of:

1. $\text{root-distance}(u)$,
2. the NCA label of u generated by Lemma 2.1,
3. the inorder number of the node u' corresponding to u in $\mathcal{C}(T)$ (so that given $u, v \in T$ we can determine which node dominates the other).

Now, suppose that u is associated with $u' \in \mathcal{C}(T)$ and v with $v' \in \mathcal{C}(T)$. We use (2) to determine $j = \text{lightdepth}(u, v) + 1$. We assume that the inorder number of u' is smaller than that of v' (using (3) we can verify this and swap u with v otherwise). Thus, u dominates v , which implies that $\text{root-distance}(\text{NCA}(u, v)) = \sum_{i=1}^j \mathbf{d}(\ell_i(u)) - 1$. Recall from Section 2 that $\text{root-distance}(\text{NCA}(u, v))$ together with $\text{root-distance}(u)$ and $\text{root-distance}(v)$ suffice to compute $\mathbf{d}(u, v)$. \square

3.2 Modified Distance Arrays

The main challenge remaining, is how to efficiently encode $D(u)$ for an arbitrary node u . This can clearly be done using $\Theta(\log^2 n)$ bits. By using properties of the heavy path decomposition, Alstrup et al. [8] gave a more precise bound of: $\sum_{i=1}^k \log \mathbf{d}(\ell_i(u)) = \sum_{i=1}^{\log n} \log(n/2^i) = 1/2 \log^2 n + O(\log n)$. In their description, sums of the suffixes of $D(u)$ are stored instead of $D(u)$ itself, but this is essentially the same. Furthermore, distance arrays must be made self-delimiting by adding an additional $O(\log n \log \log n)$ bits, so we get an overall space bound of $1/2 \log^2 n + O(\log n \log \log n)$.

In this section, we present an improved method and analysis for encoding the distance arrays. We show that our encoding uses less space, but in the process we lose the ability to compute the sum $\sum_{i=1}^j \mathbf{d}(\ell_i(u))$, which is used to answer the query. However, in Section 3.3 we show that in fact a query can still be answered by adding only a small amount of auxiliary information. Our *modified distance array* $\hat{D}(u)$ will have the following key property, which is weaker than that of the original distance array:

Property 3.2. *Given the modified distance arrays $\hat{D}(u)$ and $\hat{D}(v)$ for leaves $u, v \in T$ such that u dominates v , we can compute the value $\mathbf{d}(\ell_j(u))$ where $j = \text{lightdepth}(u, v) + 1$.*

At a high level, the main idea behind the modified distance array is that, to reduce the number of bits stored for each distance $d(\ell_1(u)), \dots, d(\ell_k(u))$ at node u , we potentially push some of the bits to labels of nodes dominated by u . This is acceptable if our goal is to satisfy Property 3.2 since we need only compute $d(\ell_i(u))$ if the other queried node v is dominated by u . An important observation for the analysis later is that if $\ell_i(u)$ is *exceptional*, we need not store $d(\ell_i(u))$ at all in order to satisfy Property 3.2. The modified distance array consists of two parts:

1. a list of *truncated distances* $\hat{d}(\ell_1(u)), \dots, \hat{d}(\ell_k(u))$;
2. a list of *accumulators* $\mathbf{a}(\ell_1(u)), \dots, \mathbf{a}(\ell_k(u))$.

Accumulator $\mathbf{a}(\ell_i(u))$ will potentially (but not necessarily) store some of the bits of the distances $d(\ell_i(v))$ where v is a node that dominates u , and $i = \text{lightdepth}(u, v) + 1$.

The construction of the labels is recursive: Consider the heavy path P extending from the root of T . Let n_1, \dots, n_{m+1} be the sizes of the subtrees T_1, \dots, T_{m+1} hanging from P via light edges e_1, \dots, e_{m+1} , where e_{m+1} is the *exceptional* edge. The edges e_1, \dots, e_{m+1} are ordered according to their left-to-right order in the collapsed tree, and we use w_1, \dots, w_m to denote the nodes in P from which e_1, \dots, e_m branch (e_{m+1} also branches from w_m). See Figure 3. We use n'_1, \dots, n'_m to denote the sizes of the subtrees T'_1, \dots, T'_m rooted at nodes w_1, \dots, w_m . For consistency, n'_{m+1} denotes the size of $T'_{m+1} = T_{m+1}$. Note that, for an arbitrary node $u \in T_i$ we have that $\ell_1(u) = e_i$.




Figure 3: A heavy path P and the subtrees T_1, \dots, T_{m+1} and T'_1, \dots, T'_m . The edge e_5 is exceptional.

For an arbitrary node $u \in T_i$ where $i \in [1, m]$, we assume that we have some encoding of its modified distance array excluding the encoding of $d(e_i) = d(\ell_1(u))$ using $(\log^2 n_i + \log n_i \log q)/4$ bits, where q is a parameter to be fixed later. We call this encoding the *recursive problem*, and the problem of encoding $d(\ell_1(u))$ the *top-level problem*. Recall that if $i = m + 1$ (i.e., $u \in T_{m+1}$) we need not encode the distance $d(e_{m+1})$, since that edge is *exceptional*.

We analyze the space of the top-level problem for T'_i for $i = m, m - 1, \dots, 1$ (i.e., from bottom to top), bounding the overall label size in terms of n'_i . The goal of each iteration is to produce labels for T'_i of size $(\log^2 n'_i + \log n'_i \log q)/4$. Consider the labels generated in the recursive problem for nodes in T_i and in the previous iteration for T'_{i+1} (or, if $i = m$, in the recursive problem for T_{m+1}). The following two lemmas show how many bits we can spend to generate the labels for nodes in T'_i . Note that these lemmas ignore the cost of making the

encoding self-delimiting, as well the fact that we must take the ceiling of the bound because we cannot store a fraction of a bit. We handle these issues later.

Lemma 3.3 (Slack Lemma). *Assume that the recursive problem for nodes in T_i can be solved by storing an encoding of size $(\log^2 n_i + \log n_i \log q)/4$ bits for some parameter q . If $n_i = p \cdot n'_i$ and $p \geq 1/q$ then we can spend additional $1/2 \log(1/p) \log n'_i$ bits on the top-level problem for nodes in T_i to obtain an encoding of size $(\log^2 n'_i + \log n'_i \log q)/4$ bits.*

Proof. To prove the lemma it is enough to calculate the difference between the size of the final encoding and the encoding for the recursive problem:

$$\begin{aligned}
&= (\log^2 n'_i + \log q \log n'_i)/4 - (\log^2(p \cdot n'_i) + \log q \log(p \cdot n'_i))/4 \\
&= (\log^2 n'_i + \log q \log n'_i - (\log p + \log n'_i)^2 - \log q \log p - \log q \log n'_i)/4 \\
&= (\log^2 n'_i + \log q \log n'_i - \log^2 p - 2 \log p \log n'_i - \log^2 n'_i - \log q \log p - \log q \log n'_i)/4 \\
&= (-\log^2 p - 2 \log p \log n'_i - \log q \log p)/4 \\
&= (2 \log(1/p) \log n'_i + \log q \log(1/p) - \log^2(1/p))/4 \\
&= (2 \log(1/p) \log n'_i + \log(1/p)(\log q - \log(1/p)))/4 \\
&\geq 1/2 \log(1/p) \log n'_i. \quad \square
\end{aligned}$$

Additionally, we have the following:

Lemma 3.4 (Thin Lemma). *Assume that the recursive problem for nodes in T_i can be solved by storing an encoding of size $(\log^2 n_i + \log n_i \log q)/4$ bits for some parameter $q \geq 2$. If $n_i = p \cdot n'_i$ and $p \leq 1/2^8$ then we can spend additional $2 \log n'_i$ bits on the top-level problem for nodes in T_i to obtain an encoding of size $(\log^2 n'_i + \log n'_i \log q)/4$ bits.*

Proof. Similarly as in the proof of Lemma 3.3, we calculate the difference:

$$\begin{aligned}
&= (\log^2 n'_i + \log q \log n'_i)/4 - (\log^2(p \cdot n'_i) + \log q \log(p \cdot n'_i))/4 \\
&= (2 \log(1/p) \log n'_i + \log(1/p)(\log q - \log(1/p)))/4.
\end{aligned}$$

Now, assuming that $2 \log n'_i$ is larger than the difference and using that $p \geq 1/n'_i$ we obtain:

$$\begin{aligned}
2 \log n'_i &> 1/2 \log(1/p) \log n'_i + 1/4 \log(1/p) \log q - 1/4 \log^2(1/p) \\
&\geq 1/2 \log(1/p) \log n'_i - 1/4 \log^2(1/p) \\
&\geq 1/2 \log(1/p) \log n'_i - 1/4 \log(1/p) \log n'_i \\
&= 1/4 \log(1/p) \log n'_i
\end{aligned}$$

so, after dividing by $\log n'_i$, $8 > \log(1/p)$ and $p > 1/2^8$. Hence for $p \leq 1/2^8$ we can indeed use $2 \log n'_i$ additional bits. \square

We call T_i *thin* if $n_i \leq n'_i/2^8$, and *fat* otherwise. We observe that, by the definition of the heavy path decomposition, $\log n \leq \log(2n'_i) \leq 2 \log n'_i$. Thus, an immediate consequence of Lemma 3.4 is that if T_i is thin, then we can afford to store $\mathbf{d}(e_i)$ explicitly as $\hat{\mathbf{d}}(e_i)$, without having to push any bits to the accumulators of nodes in T_{i+1}, \dots, T_{m+1} . However, if T_i is fat, Lemma 3.3 indicates that we do not have enough *slack* to store all the bits of $\mathbf{d}(e_i)$. Instead, we store as many bits as the slack allows (rounding up to the nearest bit) in the labels of nodes u in T_i . We then append all the remaining bits to the accumulators $\mathbf{a}(\ell_i(v))$ of all nodes $v \in \bigcup_{j=i+1}^{m+1} T_j$ (i.e., nodes dominated by u).

Because T_i is fat, by the slack lemma for nodes in T_i we have slack $1/2 \log(n'_i/n_i) \log n'_i$ (the assumption that T_i is fat allows us to adjust the constant q). On the other hand, using the same calculations as in the slack lemma, the nodes in T'_{i+1} have slack $1/2 \log(n'_i/n'_{i+1}) \log n'_i$: note that

the size of T'_{i+1} is larger than $n'_i/2^8$ by the properties of the heavy path decomposition, as either $i < m$ and $n'_{i+1} \geq n/2$, or $i = m$ and then $n_{i+1} \geq n_i$ so $n'_{i+1} \geq n/4 - 1/2$. Since $n'_i > n_i + n'_{i+1}$, we have that the sum $1/2(\log(n'_i/n_i) + \log(n'_i/n'_{i+1})) \log n'_i$ can be lower bounded by the minimum of $1/2(\log(1+x) + \log(1+x^{-1})) \log n_i$ for $x \in (0, \infty)$. Thus, the slack is at least $\log n'_i$ bits in total. However, the distance $d(\ell(u)_1)$ occupies $\log n$ bits, rather than $\log n'_i$. As before, we can use the properties of the heavy path decomposition to bound $\log n \leq \log(2n'_i) = 1 + \log n'_i$. Thus, $d(\ell_1(u))$ occupies one extra bit more than we have accounted for with the slack. We store this extra bit in the truncated distance $\hat{d}(\ell_1(u))$. Therefore, the truncated distance $\hat{d}(\ell_1(u))$ consists of the most significant $\lceil 1/2 \log(n'_i/n_i) \log n'_i \rceil + 1$ bits of $d(\ell_1(u))$. The remaining least significant $\lfloor 1/2 \log(n'_i/n'_{i+1}) \log n'_i \rfloor$ bits are concatenated to the accumulators of the nodes dominated by u in T_{i+1}, \dots, T_{m+1} .

For each entry in the modified distance array for a node u , we are pushing at most two extra bits beyond those accounted for in the slack lemma. Thus, this works out to an additional $O(\log n)$ bits in total, per label. We make both parts of the modified distance array (the accumulators and truncated distances) self-delimiting, and also record, for each truncated distance, the number of bits pushed to the accumulators of dominated nodes. Overall, we end up with the following:

Lemma 3.5. *The modified distance array $\hat{D}(u)$ occupies at most $1/4 \log^2 n + O(\log n \log \log n)$ bits.*

It remains to show that these modified distance arrays satisfy Property 3.2. To see this, consider the modified distance array for u and v , where $j = \text{lightdepth}(u, v) + 1$, and u dominates v . We have stored the number of bits that were pushed to the accumulator $\mathbf{a}(\ell_j(v))$ explicitly. The starting position of this contiguous range of bits can be found by noticing that the accumulator $\mathbf{a}(\ell_j(u))$ is a suffix of $\mathbf{a}(\ell_j(v))$, since all nodes that dominate u also dominate v . Hence, knowing the length of the accumulator $\mathbf{a}(\ell_j(u))$ allows us to determine the starting position and, together with the explicitly stored number of pushed bits, recover the bits themselves. By combining them with $\hat{d}(\ell_j(u))$ we can reconstruct $d(\ell_j(u))$.

We have therefore satisfied Property 3.2. It remains to show why this is enough for a distance query. In Section 3.3 we show that it is, with only additional lower order terms to the label size.

3.3 Wrapping Up the Proof of Theorem 1.1

In order to prove Theorem 1.1 we need to show how to answer a distance query without inflating the space of Lemma 3.5 by more than lower order terms.

Let u be some node contained in the heavy path mapped to $u' \in \mathcal{C}(T)$, and consider the path from u' to the root of $\mathcal{C}(T)$. We partition this path into $B = \sqrt{\log n}$ fragments. The first fragment is the prefix starting at the root, denoted $f_0(u)$, and terminating at the first node $f_1(u)$ such that the subtree rooted at $\text{head}(f_1(u))$ has size at most $n/2^B$. The i -th fragment is defined recursively from $f_{i-1}(u)$, ending at a node $f_i(u)$ such that the subtree rooted at $\text{head}(f_i(u))$ has size at most $n/2^{iB}$, for $i \in [1, h]$, where $h = O(\sqrt{\log n})$. We explicitly store the distances $d(f_i(u), \text{root}(T))$ for each $i \in [1, h]$ as the *fragment distance array* $F(u)$.

Next, consider a light edge e in $\mathcal{C}(T)$ that branches from the heavy path corresponding to u' to the heavy path corresponding to v' . Recall that, in bounding the number of bits for the modified distance arrays, we used the fact that if the subtree rooted at $\text{head}(u')$ has size n , then the distance, $r = d(\text{head}(u'), \text{head}(v'))$, associated with the light edge e is bounded by n . Instead of recording this distance r , for each node u that stores r we instead record the distance $r' = d(\text{head}(f_j(u)), \text{head}(v'))$, where j is the largest index such that the subtree rooted at $\text{head}(f_j(u))$ contains node $\text{head}(v')$.

Obviously, r' requires more bits to store than r , $O(\sqrt{\log n})$ additional bits to be precise. However, since there are at most $\log n$ truncated distances in $\hat{D}(u)$, we can afford to inflate each of these by $O(\sqrt{\log n})$ bits. This only increases the lower order space term to $O(\log^{1.5} n)$

bits. Furthermore, for each truncated distance, we can also afford to store the corresponding index j from the fragment array using $O(\log n \log \log n)$ extra bits. Thus, since Property 3.2 still holds after expanding the truncated distances, we can now recover r' and read $d(f_j(u), \text{root}(T))$ from $F(u)$. These two values sum to $d(\text{head}(u'), \text{root}(T))$, which is exactly what we wanted to compute with distance arrays.

The proof of Theorem 1.1 follows from the above $1/4 \log^2 n + o(\log^2 n)$ -bit labeling scheme and the fact that we only need to label leaves and can assume T is binary (see Section 2).

3.4 Query Time Analysis

Up until now we have not discussed how long it takes to compute the distance given two labels for nodes u and v . Let us summarize the steps that are required to answer a query:

1. Ensure u dominates v , and swap them if that is not the case. This can be done by examining the inorder number for u and v , which are explicitly stored; Lemma 3.1 item (3).
2. Extract the explicitly stored distances of u and v from the root; Lemma 3.1 item (1).
3. Compute the index $j = \text{lightdepth}(u, v) + 1$, this is done using the explicitly stored NCA encoding; Lemma 3.1 item (2).
4. Extract the truncated distance $\hat{d}(\ell_j(u))$ from array $\hat{D}(u)$. Note that $\hat{D}(u)$ contains $O(\log n)$ values, and has length $O(\log^2 n)$ bits.
5. Extract the accumulator values $\mathbf{a}(\ell_j(u))$ from array $\hat{D}(u)$ and $\mathbf{a}(\ell_j(v))$ from array $\hat{D}(v)$.
6. Extract explicitly stored lengths of accumulator values $\mathbf{a}(\ell_j(u))$ and $\mathbf{a}(\ell_j(v))$. Note that there are $O(\log n)$ explicitly stored lengths, and these lengths occupy $O(\log n \log \log n)$ bits.
7. Use bitwise arithmetic to extract the relevant bits of $\mathbf{a}(\ell_j(v))$ which are then concatenated with $\hat{d}(\ell_j(u))$. This can be done with a constant number of shifts, bitwise and/or operations, and subtractions.
8. Extract the fragment number for j , as well as the fragment distance from array $F(u)$. There are $O(\log n)$ fragment numbers, occupying a total of $O(\log n \log \log n)$ bits, and a total of $O(\sqrt{\log n})$ fragment distances, occupying a total of $O(\log^{1.5} n)$ bits.
9. Compute the overall distance using addition and subtraction.

With the exception of accessing the values stored in the various arrays just mentioned, all steps clearly take constant time. It remains to show how to access each array element in constant time (without increasing the space bound by more than a lower order term). First, we explicitly store the offsets of each of the (constant number of) data structures mentioned above (arrays, individual values, and the NCA labeling) for each label in a header, which is encoded using Elias δ codes in order to be self-delimiting. This header occupies at most $O(\log n)$ bits, and provides constant time access to each data structure. Next we discuss how to access the array elements in constant time. Earlier, we mentioned that we used Elias δ codes to delimit each array element and then concatenate their encodings. Now, for each array that occupies x bits in total and stores y elements, let $p_1 < p_2 < \dots < p_y$ be the positions of the first bit of the encoding of each element in the concatenation. We apply Lemma 2.2 to this sequence. This takes $O(x \cdot \max\{1, \log \frac{y}{x}\})$ and allows us to calculate the first and the last bit of the encoding of any element in constant time. For each of our arrays, $x = O(\log n)$ and $y = O(\log^2 n)$, so storing the sequences increases the total space by only $O(\log n \log \log n)$ bits. Since there is a constant number of arrays, we can afford to mark the location of their corresponding sequences in the header using $O(\log n)$ bits. Then, each array access can be performed in constant time.

3.5 Lower Bound for Level-Ancestor Labeling

In this section we prove Theorem 1.2. The main idea of the proof is to show a lower bound for *parent labeling*, where the goal is to assign a distinct label to every $u \in T$ so that given the label of u we can return the label of its parent (or a special value \perp denoting that $u = \text{root}(T)$). This is clearly a special case of the level-ancestor problem. The lower bound is obtained by showing a correspondence between the parent labeling and the following *universal tree* problem: what is the size of the smallest rooted tree containing any rooted tree on n nodes as a subtree? The connection between these two problems is captured by the following lemma.

Lemma 3.6. *If there exists a parent labeling scheme on trees of size n that produces labels of size at most $S(n)$, then there exists a universal rooted tree containing all rooted trees on up to n nodes as subtrees of size $O(2^{S(n)})$.*

Proof. The proof is by construction. Let V be the set of all possible labels generated by the labeling scheme, and E be the directed edges between these labels defined as follows: if a label u is assigned to a node of some tree, and $v \neq \perp$ is the label returned by the scheme for u , then (u, v) belongs to E . Note that v is determined solely from the bits of u , hence the graph $G = (V, E)$ consists of one or more directed cycles. See Figure 4 (left) for an example of such a graph. It is clear that G must contain any tree T on up to n nodes as a subgraph, since the labeling scheme works for all trees on n nodes or less. G is not necessarily a tree itself, but we now describe a general procedure that converts G into a new graph $G' = (V', E')$ that itself is a rooted tree, and is such that that $|V'| \leq 2|V| + 1$.

Each weakly connected component of G is either already a tree, or contains a cycle. In the latter case, we arbitrarily remove an edge (u, v) from the cycle (in the figure the chosen edge is intersected by the dashed line). After deleting (u, v) we duplicate the entire weakly connected component, and add a new edge (u, v') where v' is the duplicate of v . After doing this for each weakly connected component, we have increased the number of vertices to at most $2|V|$, and the resultant graph G' is a forest of rooted trees. We add a single global root to make G' a rooted tree. The total number of nodes in G' is hence at most $2|V| + 1$.

Since G was a universal graph for rooted trees on n nodes, any rooted tree not containing the deleted edge clearly appears as a subgraph in G' . Moreover, for any rooted tree T containing (u, v) , there exists some subpath of the cycle in G which was in T . Since we duplicated each node in the cycle, it is clear that any such subpath also exists in G' (together with any trees rooted at nodes in the subpath), thus, T appears as a subtree in G' .

The final detail is to consider the maximum length label output by the labeling scheme, which consists of $S(n)$ bits. Hence, there are at most $2^{S(n)}$ nodes in G and therefore at most $O(2^{S(n)})$ nodes in G' . \square

Equipped with the previous lemma, we immediately get a lower bound on $S(n)$, provided we have a lower bound on the number of nodes in such a rooted universal tree. Goldberg and Lifschitz [22] have proved very accurate bounds on the number of nodes in such rooted universal trees (see [14] for the bound as we state it):

Lemma 3.7 ([14, 22]). *The smallest rooted tree containing all rooted trees on up to n nodes as subtrees has size $n^{(\log n - 2 \log \log n + O(1))/2}$.*

By combining Lemmas 3.6 and 3.7, we obtain Theorem 1.2.

3.6 Effective Level-Ancestor Scheme

While Alstrup et al. [8] describe their scheme in terms of labeling for distance queries, in fact it is not difficult to tweak it to obtain a scheme for level-ancestor queries. We describe the necessary




Figure 4: Converting a weakly connected component to a rooted tree G' by duplicating the path at the dotted line.

modifications to obtain a scheme for parent queries. repeatedly moving to the parent as long as necessary.

The labeling consists of three parts. For a node u on a heavy path P we store:

1. $d(u, \text{root}(T))$,
2. the $O(\log n)$ label generated by Lemma 2.1 applied on $\text{head}(P)$,
3. the array $D(u)$ and, additionally, $d(u, \text{head}(P))$. (This is differently phrased but essentially equivalent to what the original labeling stores.)

The labels in the NCA scheme are required to be distinct, so the labels of nodes belonging to different heavy paths are distinct. For two nodes on the same heavy path, storing $d(u, \text{head}(P))$ explicitly ensures that their labels are not the same. Each label consists of $1/2 \log^2 n + O(\log n)$ bits, because of the bound on the encoding of $D(u)$. We need to argue that given the label of $u \neq \perp$ we can construct the label of its parent.

The NCA labeling scheme from Lemma 2.1 has the property that the label of every node u is a concatenation of heavy and light labels $h_0 \ell_1 h_1 \ell_2 \cdots \ell_k h_k$. These labels uniquely determine the path from the root to u : h_0 encodes how far along the heavy path starting at the root we should continue. Then, either $k = 0$ and u in fact lies on the heavy path starting at the root, or ℓ_1 encodes which light edge outgoing from the current node should be followed. Finally, $h_1 \ell_2 \cdots \ell_k h_k$ recursively encodes the remaining part of the path to u in the subtree hanging off the heavy path starting at the root. It is not necessarily true that given the NCA label of a node u we can determine the NCA label of its parent. However, by truncating the NCA label of u we can construct the NCA label of the parent of $\text{head}(P)$.

Given the label of u , we construct the label of its parent u' as follows. $d(u, \text{root}(T))$ needs to be decreased by 1. Then we inspect $d(u, \text{head}(P))$. If $u \neq \text{head}(P)$, we decrease $d(u, \text{head}(P))$ by 1 and are done. Otherwise, we use the NCA label of u to determine the label of its parent as explained above. Let P' be the heavy path of u' . The last element of the array D is $d(u, \text{head}(P'))$, so by subtracting 1 we obtain $d(u', \text{head}(P'))$. Finally, we remove the last element of D .

4 k -Distance Labeling

Recall that in k -distance labeling, given the labels of u and v we need to decide if the length of the u -to- v path is at most k , and if so return it. In this section we prove Theorem 1.3.

4.1 Lower Bound for Small k

We define a family of trees and show that labeling the leaves of all trees in that family for k -distance queries requires $\log n + \Omega(k \cdot \log \frac{\log n}{k \log k})$ bits.

An \vec{x} -regular tree, where $\vec{x} = (x_1, \dots, x_k) \in \mathbb{N}^k$, is a rooted tree of height k where all depth- i nodes have the same degree x_{i+1} . An (\vec{x}, h, d) -regular tree, where $(x_1, \dots, x_k) = \vec{x} \in [h]^k$, is a \vec{y} -regular tree with $\vec{y} = (d^{x_1}, d^{h-x_1}, d^{x_2}, d^{h-x_2}, \dots, d^{x_k}, d^{h-x_k})$. The total number of leaves in such a tree is $d^{k \cdot h}$. See Figure 5 for an example.




Figure 5: (\vec{x}, d, h) -regular tree with $\vec{x} = (1, 2)$ and $d = h = 2$

We consider (\vec{x}, h, d) -regular trees for some parameters h and d to be chosen later. Consider a labeling scheme that assigns a label to every leaf of such a tree for $2k$ -distance queries. The following lemma shows that a (\vec{x}, h, d) -regular tree and a (\vec{y}, h, d) -regular tree cannot share many identical labels. More formally, let $\text{common}(\vec{x}, \vec{y})$ denote the maximum number of labels that can be used in both trees. The following is an upper bound on the sum of $\text{common}(\vec{x}, \vec{y})$.

Lemma 4.1. $\sum_{\vec{x}, \vec{y} \in [h]^k} \text{common}(\vec{x}, \vec{y}) \leq \left(h \cdot d^h \left(1 + \frac{2}{d-1} \right) \right)^k$.

Proof. We first prove that

$$\text{common}(\vec{x}, \vec{y}) \leq \prod_{i=1}^k d^{\min\{x_i, y_i\}} d^{h - \max\{x_i, y_i\}}.$$

By asking all $2k$ -distance queries between a specified subset S of leaves of the (\vec{x}, h, d) -regular tree we can recover the shape of the subtree induced by S . Hence, if two trees share $\text{common}(\vec{x}, \vec{y})$ labels, then they must have a common isomorphic subtree on $\text{common}(\vec{x}, \vec{y})$ leaves. To bound the maximum number of leaves in such a subtree, observe that the degree of a node at depth $2i - 2$ is at most $\min\{d^{x_i}, d^{y_i}\}$, and the degree of a node at depth $2i - 1$ is at most $\min\{d^{h-x_i}, d^{h-y_i}\}$. The maximum number of shared labels is hence the product of all these quantities over $i = 1, \dots, k$. We conclude that that $\text{common}(\vec{x}, \vec{y}) \leq \prod_{i=1}^k d^{\min\{x_i, y_i\}} d^{h - \max\{x_i, y_i\}}$. It then follows that

$$\begin{aligned} \sum_{\vec{x}, \vec{y} \in [h]^k} \text{common}(\vec{x}, \vec{y}) &= \sum_{\vec{x}, \vec{y} \in [h]^k} \prod_{i=1}^k d^{\min\{x_i, y_i\}} d^{h - \max\{x_i, y_i\}} \\ &= \prod_{i=1}^k \sum_{1 \leq x, y \leq h} d^{\min\{x, y\}} d^{h - \max\{x, y\}} \\ &= \prod_{i=1}^k \left(h \cdot d^h + 2 \sum_{x < y} d^x d^{h-y} \right) \\ &= \prod_{i=1}^k \left(h \cdot d^h + 2 \sum_{x=1}^{h-1} d^x \sum_{y=0}^{h-x-1} d^y \right) \\ &= \prod_{i=1}^k \left(h \cdot d^h + 2 \sum_{x=1}^{h-1} d^x \frac{d^{h-x} - 1}{d - 1} \right) \end{aligned}$$

$$\begin{aligned} &\leq \prod_{i=1}^k \left(h \cdot d^h + 2 \cdot h \frac{d^h}{d-1} \right) \\ &= \left(h \cdot d^h \left(1 + \frac{2}{d-1} \right) \right)^k. \square \end{aligned}$$

Since the total number of leaves in the (\vec{x}, h, d) -regular trees family is $d^{k \cdot h} \cdot h^k$, the number of distinct labels required to label them is thus at least

$$d^{k \cdot h} \cdot h^k - \sum_{\vec{x} < \vec{y}} \text{common}(\vec{x}, \vec{y}) = d^{k \cdot h} \cdot h^k - \frac{1}{2} \sum_{\vec{x} \neq \vec{y}} \text{common}(\vec{x}, \vec{y}) = \frac{3}{2} \cdot d^{k \cdot h} \cdot h^k - \frac{1}{2} \sum_{\vec{x}, \vec{y} \in [h]^k} \text{common}(\vec{x}, \vec{y}).$$

Now we set $d = 2k + 1$, and since $(1 + \frac{1}{k}) \leq e^{\frac{1}{k}}$ we have from Lemma 4.1 that $\sum_{\vec{x}, \vec{y}} \text{common}(\vec{x}, \vec{y}) \leq e \cdot d^{k \cdot h} \cdot h^k$, so the number of unique labels is at least: $(3/2 - e/2) \cdot d^{k \cdot h} \cdot h^k > 0.1 \cdot d^{k \cdot h} \cdot h^k$. Setting $n = d^{k \cdot h}$ this becomes $0.1 \cdot h^k \cdot n$, making the number of required bits at least:

$$\log n + k \log h - O(1) = \log n + k \log \frac{\log n}{k \log d} - O(1) = \log n + \Omega(k \cdot \log \frac{\log n}{k \log k}).$$

Note that for the above calculation to make sense, we need that $d^k \leq n$.

4.2 Lower Bound for Large k

The lower bound from Section 4.1 is not meaningful for large values of $k \in [\log n, n]$. In this section we show that the lower bound of Gavaille et al. [21] for general distance queries, can be translated into a lower bound of $\Omega(\log n \cdot \log(k/\log n))$ for k -distance queries.

The lower bound uses the family of (h, M) -trees (see Section 2). Recall that every edge of an (h, M) -tree has a weight from $[0, M]$. It is easy to verify that the number of nodes in such a tree is $3 \cdot 2^h - 2$, hence the distance between any two leaves is no more than $2hM$.

If $M \leq k/(2h)$ then, because the distance between any two leaves in the tree is at most $2hM \leq k$, any labeling of the leaves for k -distance can be used for general distance labeling. By Lemma 2.3, such a labeling scheme would require labels of at least $h/2 \cdot \log M$ bits. We set $h = \log \sqrt{n/3}$ and $M = \min\{k/2h, 2^h\}$. Then, by subdividing the edges of an (h, M) -tree we obtain an unweighted tree on at most n nodes. Labeling the leaves of such a tree for k -distance can be used for general distance labeling of the (h, M) -tree, so we obtain the following lower bounds:

- (1) if $\frac{k}{2h} \leq 2^h$, the number of required bits is $\frac{h}{2} \cdot \log \frac{k}{2h} = \Omega(\log n \cdot \log \frac{k}{\log n})$;
- (2) if $\frac{k}{2h} > 2^h$, the number of required bits is $\frac{h}{2} \cdot h = \Omega(\log^2 n)$, so $\Omega(\log n \cdot \log \frac{k}{\log n})$ for $k \leq n$.

4.3 Upper Bound

In this section we present our improved upper bound for k -distance labeling. We build upon the ideas of Alstrup, Bille, and Rauhe [5], who presented an $\log n + O(k^2 \log(k \log n))$ bits labeling scheme. As a preliminary step, we will show an $O(\log k \cdot \log n)$ bits scheme for $k \geq \log n$, and then move to the more complicated $\log n + O(k \log(\log n/k))$ bits scheme for $k < \log n$.

Consider the heavy path decomposition of T . We define the *light range* of u , denoted L_u , to contain the preorder number of all nodes in T_u if u has no heavy child, and all nodes in $T_u \setminus T_{\text{heavy}(u)}$ otherwise. We say that v is a *significant ancestor* of u if $\text{pre}(u) \in L_v$. For example, in Figure 1 v is a significant ancestor of u since the light range of v is $L_v = [5, 23)$. The number of significant ancestors of u is equal to $\text{lightdepth}(u) = O(\log n)$. The nearest common significant ancestor of u and v , denoted $\text{NCSA}(u, v)$, is w such that $\text{pre}(w)$ is as large as possible and w is

a significant ancestor of both u and v . In other words, w is the first significant ancestor on the path from u to the root, which is also a significant ancestor of v . The heavy path P such that $\text{head}(P)$ is a child of $\text{NCSA}(u, v)$ is called the nearest common heavy path of u and v and denoted $\text{NCH}(u, v)$. When there is no common significant ancestor for u and v we set $\text{NCSA}(u, v)$ to nil and $\text{NCH}(u, v)$ to be the heavy path starting at the root.

Let the significant ancestors of u and v on $\text{NCH}(u, v)$ be u' and v' , respectively. Then $d(u, v) = d(u, u') + d(u', v') + d(v, v')$. Computing $d(u, v)$ consists of two steps:

1. identifying $\text{NCH}(u, v)$, u' and v' , and computing $d(u, u')$ and $d(v, v')$,
2. computing $d(u', v')$.

We describe these steps separately, and then describe how to implement them in constant time.

Identifying $\text{NCH}(u, v)$. For an integer range $A = [a, b] \subset [1, n]$ we define its identifier $\text{id}(A)$ by considering a binary trie representing all words of length $\lceil \log n \rceil$. The label of a node u in the trie is the concatenation of the labels of the edges on the path from the root to u . Every integer $x \in [1, n]$ corresponds to a leaf u in the trie, such that the label of u is the binary expansion of x . Then, $\text{NCA}(a, b)$ is the nearest common ancestor of the leaves corresponding to a and b in the trie, $\text{height}(A)$ is the height of the subtree rooted at $\text{NCA}(a, b)$, and finally $\text{id}(A)$ is the label of $\text{NCA}(a, b)$.

Observation 4.2. For any range A :

1. $\text{id}(A)$ can be computed given $\text{height}(A)$ and any $x \in A$,
2. $A \cap B = \emptyset \implies \text{id}(B) \neq \text{id}(A)$.

Alstrup, Bille, and Rauhe [5] use the notion of *significant preorder numbers*. We replace it with our notion of range identifier, that has very similar properties, yet is somewhat easier to operate on (and hence we are able to achieve a better query time). For any node $u \in T$, let $\text{id}(u) = (\text{id}(L_u), \text{lightdepth}(u))$.

Lemma 4.3. For any nodes $u, v \in T$, if $u \neq v$ then $\text{id}(u) \neq \text{id}(v)$.

Proof. If $\text{lightdepth}(u) \neq \text{lightdepth}(v)$ then we are done. Otherwise, L_v and L_u are disjoint, so by Observation 4.2.2 $\text{id}(L_v) \neq \text{id}(L_u)$ and we are also done. \square

Consider a node $u \in T$ and let $u = u_0, u_1, u_2, \dots$ be all of its significant ancestors in the order in which they appear on the path from u to $\text{root}(T)$. Let u_r be the last of these ancestors such that $d(u, u_r) \leq k$. We call u_r the top significant ancestor of u . The label of u consists of $\text{pre}(u)$, $\text{lightdepth}(u)$, and an encoding of $\text{height}(L_{u_i})$ for every $i = 0, 1, \dots, r$. By Observation 4.2.1 this is enough to compute $\text{id}(u_i)$ for every $i = 0, 1, \dots, r$. Consequently, given the labels of u and v , we can either detect that the distance from u or v to $\text{NCA}(u, v)$ exceeds k , or calculate $\text{lightdepth}(\text{NCSA}(u, v))$.

To encode $\text{height}(L_{v_i})$ for every $i = 0, 1, \dots, r$, we observe that $L_{v_i} \subseteq L_{v_{i+1}}$ and that $r \leq \min\{\log n, k\}$. Hence, we need to encode a non-decreasing sequence of $\min\{\log n, k\}$ numbers from $[0, \log n]$. By Lemma 2.2, for $k < \log n$ this can be done using $O(k \log(\log n/k))$ bits and for $k \geq \log n$ using $O(\log n)$ bits, and allows us to calculate $\text{lightdepth}(\text{NCSA}(u, v))$ or detect that $d(u, v) > k$.

We encode in the label of u the distance from u to u_i for every $i = 0, 1, \dots, r-1$. Because $0 = d(u, u_0) < d(u, u_1) < \dots < d(u, u_{r-1}) \leq k$ we need to encode an increasing sequence of $\min\{\log n, k\}$ numbers from the range $[0, k]$. By Lemma 2.2, if $k < \log n$ this can be done using $O(k)$ bits and if $k \geq \log n$ using $O(\log n \cdot \log(k/\log n))$ bits. Then, after having found $\text{lightdepth}(\text{NCSA}(u, v))$ we can compute $d(u, u')$ and $d(v, v')$.




Figure 6: $w = \text{NCSA}(u, v)$, u' and v' are the significant ancestors of u and v on $\text{NCH}(u, v)$, respectively. Significant ancestors are white, heavy edges are solid, and light edges are dashed.

Computing $d(u', v')$. Recall that u' and v' are the significant ancestors on the $\text{NCH}(u, v)$ of u and v , respectively. We want to compute $d(u', v')$. If u' is not the top significant ancestor of u and v' is not the top significant ancestor of v then from the distances encoded in the labels of u and v we can retrieve $d(u', \text{NCSA}(u, v))$ and $d(v', \text{NCSA}(u, v))$, and return their absolute difference as $d(u', v')$. Now consider the case that u' is the top significant ancestor of u , but v' is not the top significant ancestor of v . To deal with this case, the label of u should also encode the distance α from u' to the head of its heavy path. This distance might be very large (even up to n), so we cap it at $2k + 1$ to use only $O(\log k)$ bits. Since v' is not the top significant ancestor of v , we can retrieve $\beta = d(v', \text{NCSA}(u, v))$ as in the previous case. We know that $\beta \leq k$ because otherwise v' would be the top significant ancestor of v . Recall that $d(u', v')$ is equal to the absolute difference between $d(u', \text{NCSA}(u, v))$ and $d(v', \text{NCSA}(u, v))$. If $\alpha = 2k + 1$ then this value must exceed k , so we terminate. Otherwise, we return $|\alpha - \beta|$.

The remaining and most complicated case is when u' is the top significant ancestor of u and v' is the top significant ancestor of v . If $k > \log n$, the solution is simple, as we can afford to store the distance from the top significant ancestor to the head of its heavy path for every node (i.e., $d(u, u_{r+1})$) using $O(\log n)$ bits. The rest of this section is dedicated for solving $k \leq \log n$.

To make the further exposition more concise, we define the 2-approximation of an integer x , denoted $\lfloor x \rfloor_2$, as the largest power of 2 not exceeding x . That is, $\lfloor x \rfloor_2 = 2^{\lfloor \log x \rfloor}$. Clearly, 2-approximation is monotone, meaning that $x \leq y$ implies $\lfloor x \rfloor_2 \leq \lfloor y \rfloor_2$, and furthermore $\lfloor x \rfloor_2 < \lfloor 2x \rfloor_2$.

Lemma 4.4. *Let A, B, C be three open intervals such that $A \cap B = \emptyset$ and $A, B \subseteq C$. Then $\lfloor |C| \rfloor_2 \neq \lfloor |A| \rfloor_2$ or $\lfloor |C| \rfloor_2 \neq \lfloor |B| \rfloor_2$.*

Proof. Assume that $|B| \leq |A|$. Then $2|B| \leq |A| + |B| \leq |C|$ and by the properties of 2-approximation $\lfloor |B| \rfloor_2 < \lfloor |C| \rfloor_2$, so indeed $\lfloor |B| \rfloor_2 \neq \lfloor |C| \rfloor_2$. Symmetrically, if $|A| \leq |B|$ then $\lfloor |A| \rfloor_2 \neq \lfloor |C| \rfloor_2$. \square

The following lemma captures the essence of the k -distance scheme of Alstrup, Bille, and Rauhe [5], while being optimized so that we can obtain our improvement.

Lemma 4.5. *Consider an increasing sequence of integers $a_1 < a_2 < \dots < a_s$. Given $a_i < a_j$, $i' = i \bmod k$, $j' = j \bmod k$, and $\lfloor a_{i+t} - a_i \rfloor_2$ and $\lfloor a_j - a_{j-t} \rfloor_2$ for every $t = 1, 2, \dots, k$ we can calculate $j - i$ or determine that $j - i > k$ in constant time.*

Proof. We start by setting $t = j' - i'$. Now either $t = j - i$ or $j - i \geq k + t$. Hence we only need to distinguish between these two cases.

Consider three intervals (a_i, a_{i+t}) , (a_{j-t}, a_j) and (a_i, a_j) . If $t = j - i$ then these three intervals are equal and so are $\lfloor a_{i+t} - a_i \rfloor_2$, $\lfloor a_j - a_{j-t} \rfloor_2$ and $\lfloor a_j - a_i \rfloor_2$. Otherwise $j - i \geq k + t > 2t$, so (a_i, a_{i+t}) and (a_{j-t}, a_j) are two disjoint intervals contained in (a_i, a_j) . Hence by Lemma 4.4 either $\lfloor a_{i+t} - a_i \rfloor_2 \neq \lfloor a_j - a_i \rfloor_2$ or $\lfloor a_j - a_{j-t} \rfloor_2 \neq \lfloor a_j - a_i \rfloor_2$. Therefore after retrieving $\lfloor a_{i+t} - a_i \rfloor_2$ and $\lfloor a_j - a_{j-t} \rfloor_2$ and calculating $\lfloor a_j - a_i \rfloor_2$ we can distinguish between the two cases and either return $j - i$ or report that $j - i > k$. Notice that $\lfloor a_j - a_i \rfloor_2$ can be calculated in constant time using standard word-RAM operations. \square

We need to show that, for every heavy path, we store enough information for applying Lemma 4.5. Consider a heavy path $u_1 - u_2 - \dots - u_s$, where u_1 is the head. By the properties of the heavy path decomposition, $\text{id}(L_{u_1}) < \text{id}(L_{u_2}) < \dots < \text{id}(L_{u_s})$. The label of every $u \in T$ such that u_i is the top significant ancestor of u encodes the following:

1. $\text{id}(L_{u_i})$;
2. $\lfloor \text{id}(L_{u_{i+t}}) - \text{id}(L_{u_i}) \rfloor_2$ and $\lfloor \text{id}(L_{u_i}) - \text{id}(L_{u_{i-t}}) \rfloor_2$ for every $t = 1, 2, \dots, k$;
3. $i \bmod k$.

To encode $\text{id}(L_{u_i})$, we store $\text{height}(L_{u_i})$ using $O(\log \log n)$ bits. Encoding $\lfloor \text{id}(L_{u_{i+t}}) - \text{id}(L_{u_i}) \rfloor_2$ and $\lfloor \text{id}(L_{u_i}) - \text{id}(L_{u_{i-t}}) \rfloor_2$ for every $t = 1, 2, \dots, k$ reduces to encoding two non-decreasing sequences of k integers from $[0, \log n]$. By Lemma 2.2, such a sequence can be stored using $O(k \log \frac{\log n}{k})$ -bits. Finally, $i \bmod k$ is encoded using $O(\log k)$ bits. Notice that both $O(\log \log n)$ and $O(\log k)$ are absorbed by $O(k \log \frac{\log n}{k})$.

To conclude, given the labels of u and v , whose significant ancestors u' and v' are on $\text{NCH}(u, v) = u_1 - u_2 - \dots - u_s$ and are both the top significant ancestors, we can now calculate $d(u', v')$ or detect that it exceeds k by retrieving the necessary information from the labels of u and v and then applying Lemma 4.5. Finally, in the following section (4.4) we show that queries can be supported in constant time. The gist of the improvement in the query time is that $\text{id}(L_{u_i})$ can be obtained from $\text{pre}(u)$ by truncating the last $\text{height}(L_{u_i})$ trailing bits and setting the $\text{height}(L_{u_i})^{\text{th}}$ bit to 1.

4.4 Query Time Analysis

We now show how to implement the query in constant time. The main difficulty is in determining $\text{lightdepth}(\text{NCSA}(u, v))$ efficiently. Once it is known, from $\text{lightdepth}(u)$ and the encoding of the distances from u to its significant ancestors implemented with Lemma 2.2 we obtain $d(u, u')$ in constant time, and similarly for $d(v, v')$ (or conclude that $d(u, v)$ exceeds k). Calculating $d(u', v')$ requires invoking Lemma 4.5 while providing access to the stored non-decreasing sequences of 2-approximations with Lemma 2.2, so also takes only constant time.

Recall that the label of u contains $\text{pre}(u)$, $\text{lightdepth}(u)$, and an encoding of the sequence $\text{height}(L_{u_0}) \leq \text{height}(L_{u_1}) \leq \dots \leq \text{height}(L_{u_r})$ implemented with Lemma 2.2. Similarly, the label of v contains $\text{pre}(v)$, $\text{lightdepth}(v)$, and $\text{height}(L_{v_0}) \leq \text{height}(L_{v_1}) \leq \dots \leq \text{height}(L_{v_s})$. We want to calculate $\text{lightdepth}(\text{NCSA}(u, v))$. For now, we assume that $r = s$ and $\text{lightdepth}(u_i) = \text{lightdepth}(v_i)$ for every $i = 0, 1, \dots, r$. Then, calculating $\text{lightdepth}(\text{NCSA}(u, v))$ reduces to finding the smallest i such that $u_i = v_i$. Notice that then $u_j = v_j$ for every $j = i, i + 1, \dots, r$. If $u_j = v_j$ then clearly $\text{height}(L_{u_j}) = \text{height}(L_{v_j})$, so we start with locating the smallest i' such that $\text{height}(L_{u_j}) = \text{height}(L_{v_j})$ for every $j = i', i' + 1, \dots, r$. This can be done in constant time by computing the longest common suffix of both sequences.

Because $\text{lightdepth}(u_j) = \text{lightdepth}(v_j)$ for every $j = 0, 1, \dots, r$, it remains to find the smallest $i \geq i'$ such that $\text{id}(L_{u_j}) = \text{id}(L_{v_j})$ for every $j = i, i + 1, \dots, r$. Observe that $\text{id}(L_{u_j})$ is obtained by clearing all $\text{height}(L_{u_j})$ least significant bits of $\text{pre}(u)$ and, if $\text{height}(L_{u_j}) > 0$, setting the $\text{height}(L_{u_j})^{\text{th}}$ bit to 1, and similarly for $\text{id}(L_{v_j})$. Without loss of generality, assume that

$\text{height}(L_{u_{i'}}) = \text{height}(L_{v_{i'}}) > 0$ (if not, $i = i'$ is checked separately in constant time). We find the longest common prefix of the binary expansions of $\text{pre}(u)$ and $\text{pre}(v)$, i.e., the smallest $\ell \geq 0$ such that their binary expansions are the same after truncating the ℓ least significant bits. ℓ can be found in constant time using standard word-RAM operations $\text{MSB}(\text{pre}(u) \text{ XOR } \text{pre}(v))$. Then, for $\text{id}(L_{u_i}) = \text{id}(L_{v_i})$ to hold, we need to clear at least ℓ least significant bits of $\text{pre}(u)$ and $\text{pre}(v)$. Hence it remains to find the smallest $i \geq i'$ such that $\text{height}(L_{u_i}) = \text{height}(L_{v_i}) \geq i$. Such an i can be found in constant time with a successor query on the encoded sequence.

If $r \neq s$ or $\text{lightdepth}(u_0) \neq \text{lightdepth}(v_0)$, then essentially the same argument works, except that we need to compute the longest common prefix of suffixes instead of whole sequences.

5 Approximate Distance Labeling

In this section we prove Theorem 1.4. Recall that in $(1 + \varepsilon)$ -approximate distance labeling, given the labels of u and v we need to output some value in the interval $[\mathbf{d}(u, v), (1 + \varepsilon) \cdot \mathbf{d}(u, v)]$.

5.1 Lower bound

To show the lower bound we modify the family of (h, M) -trees such that exact distances between leaves can be inferred from their approximate distances. Thereafter, we can invoke Lemma 2.3 to establish the lower bound.

An (h, M) -tree is modified by first subdividing its edges to obtain an unweighted tree of height $h \cdot M$. The edges of this unweighted tree are then further subdivided: every edge of depth $d \geq 0$ is subdivided into $\lfloor (1 + \varepsilon)^{hM-d} \rfloor$ edges. Note that in the original (h, M) -tree all leaves are at the same distance from the root. Therefore, if the distance between two leaves is $2k$ in the original tree, it is $f(k) = 2 \sum_{i=1}^k \lfloor (1 + \varepsilon)^i \rfloor$ in the final tree. A $(1 + \varepsilon)$ -approximation of this distance belongs to the interval $[f(k), (1 + \varepsilon)f(k)]$. We next show that these intervals are disjoint, so in fact a $(1 + \varepsilon)$ -approximation of $f(k)$ is enough to infer the original distance, $2k$.

Observe that $f(k)$ is monotone, so to prove that the intervals $[f(k), (1 + \varepsilon)f(k)]$ are disjoint, it is enough to show that $(1 + \varepsilon)f(k) < f(k + 1)$, or:

$$(1 + \varepsilon) \sum_{i=1}^k \lfloor (1 + \varepsilon)^i \rfloor < \sum_{i=1}^{k+1} \lfloor (1 + \varepsilon)^i \rfloor, \text{ or equivalently}$$

$$\varepsilon \sum_{i=1}^k \lfloor (1 + \varepsilon)^i \rfloor < \lfloor (1 + \varepsilon)^{k+1} \rfloor.$$

Since $\lfloor (1 + \varepsilon)^i \rfloor < (1 + \varepsilon)^i$, it is enough to show that:

$$\varepsilon \sum_{i=1}^k (1 + \varepsilon)^i < \lfloor (1 + \varepsilon)^{k+1} \rfloor, \text{ or equivalently}$$

$$(1 + \varepsilon)^{k+1} - (1 + \varepsilon) < \lfloor (1 + \varepsilon)^{k+1} \rfloor.$$

Since $x - 1 < \lfloor x \rfloor$ is always true, we conclude that the intervals are indeed disjoint. Hence, by Lemma 2.3 we obtain that labeling the leaves of the final tree for $(1 + \varepsilon)$ -approximate distances requires $h/2 \cdot \log M$ bits. It remains to choose h and M and rephrase this bound in terms of the size of the final tree. The size of the final tree is at most

$$2 \sum_{i=0}^{h-1} 2^{h-1-i} \sum_{j=M \cdot i+1}^{M \cdot (i+1)} \lfloor (1 + \varepsilon)^j \rfloor = \sum_{i=0}^{h-1} 2^{h-i} \sum_{j=M \cdot i+1}^{M \cdot (i+1)} (1 + \varepsilon)^j$$

$$\begin{aligned}
&\leq 2^h \sum_{i=0}^{h-1} 2^{-i} (1+\varepsilon)^{M \cdot i + 1} \frac{(1+\varepsilon)^M - 1}{(1+\varepsilon) - 1} \\
&\leq 2^h \frac{1}{\varepsilon} \sum_{i=0}^{h-1} 2^{-i} (1+\varepsilon)^{M(i+1)+1} \\
&\leq 2^h \frac{1}{\varepsilon} (1+\varepsilon)^{M+1} \sum_{i=0}^{h-1} \left(\frac{(1+\varepsilon)^M}{2} \right)^i \\
&= 2^h \frac{1}{\varepsilon} (1+\varepsilon)^{M+1} \frac{\left(\frac{(1+\varepsilon)^M}{2} \right)^h - 1}{\frac{(1+\varepsilon)^M}{2} - 1} \\
&\leq \frac{2}{\varepsilon} \frac{(1+\varepsilon)^{M+1}}{(1+\varepsilon)^M - 2} (1+\varepsilon)^{M \cdot h}
\end{aligned}$$

We set $M = 2/\varepsilon$. Then, because $\varepsilon \leq 1$ and $(1+\varepsilon)^M \geq 4$, the size is at most:

$$\begin{aligned}
&\leq 2 \frac{1+\varepsilon}{\varepsilon} \frac{(1+\varepsilon)^M}{(1+\varepsilon)^M - 2} e^{2h} \\
&\leq \frac{8}{\varepsilon} e^{2h}.
\end{aligned}$$

We set $h = \log(\varepsilon \cdot n/8)/(2 \log e) = \Theta(\log(\varepsilon \cdot n))$, and obtain that labeling trees of size n for $(1+\varepsilon)$ -approximate distances requires $\Omega(\log(1/\varepsilon) \cdot \log(\varepsilon \cdot n))$ bits. Now, if $\varepsilon > 1/\sqrt{n}$ this is in fact $\Omega(\log(1/\varepsilon) \cdot \log n)$ and we are done. Otherwise ($\varepsilon \leq 1/\sqrt{n}$), we observe that a scheme with such small ε can be used for labeling a tree of size \sqrt{n} for exact distances (by subdividing every edge into \sqrt{n} edges). Such labeling requires $\Omega(\log^2(\sqrt{n})) = \Omega(\log^2 n)$ bits, which for $\varepsilon \geq 1/n$ is also $\Omega(\log(1/\varepsilon) \cdot \log n)$ as required.

5.2 Upper bound

We now describe a matching upper bound: a $(1+\varepsilon)$ -approximate distance labeling scheme with label size $O(\log(1/\varepsilon) \cdot \log n)$. Our scheme is based on the scheme of Alstrup et al. [8] whose label size is $O(1/\varepsilon \cdot \log n)$. For any node v , let v_1, \dots, v_k be the significant ancestors of v in the order they appear on the v -to-root path. Let $\lceil x \rceil_{1+\varepsilon}$ denote the smallest power of $1+\varepsilon$ larger than x . Observe that $\lceil x \rceil_{1+\varepsilon}$ is a $(1+\varepsilon)$ -approximation of x .

The label of a node v in [8] is composed of the following fields:

1. $d(v, \text{root}(T))$,
2. the $O(\log n)$ label generated by Lemma 2.1 applied on v ,
3. the sequence $\lceil d(v, v_1) \rceil_{1+\varepsilon}, \lceil d(v, v_2) \rceil_{1+\varepsilon}, \dots, \lceil d(v, v_k) \rceil_{1+\varepsilon}$.

Let $w = \text{NCA}(u, v)$. If $w = v$ or $w = u$, we can extract the exact distance from (1). Otherwise, w.l.o.g. we can find the significant ancestor v_j of v such that $v_j = w$ using (2), and then find $\lceil d(v, w) \rceil_{1+\varepsilon}$ using (3). Alstrup et al. show that:

$$d(u, v) \leq d(u, \text{root}(T)) - d(v, \text{root}(T)) + 2 \cdot \lceil d(v, w) \rceil_{1+\varepsilon} \leq (1 + 2\varepsilon) \cdot d(u, v).$$

This means we can compute a $(1+\varepsilon)$ -approximation of $d(u, v)$ by replacing ε with $\varepsilon/2$. The bottleneck for the size of the label is storing the sequence in (3). In [8], this sequence is stored using a unary encoding of the sequence $\lceil d(v, v_1) \rceil_{1+\varepsilon}, \lceil d(v, v_2) \rceil_{1+\varepsilon} - \lceil d(v, v_1) \rceil_{1+\varepsilon}, \dots, \lceil d(v, v_k) \rceil_{1+\varepsilon} - \lceil d(v, v_{k-1}) \rceil_{1+\varepsilon}$ delimited by a single bit between two consecutive values. The maximal length of

the path is at most n , so such an encoding will require $\log_{1+\varepsilon} n$ bits and additional $k \leq \log n$ bits for the delimiters. This means that the final label size is $\Theta(\log_{1+\varepsilon} n)$, or $\Theta(1/\varepsilon \cdot \log n)$ for small ε . Instead, we store the sequence using Lemma 2.2, which yields a label of size $O(\log(1/\varepsilon) \cdot \log n)$ bits and a constant query time.

References

- [1] S. Abiteboul, S. Alstrup, H. Kaplan, T. Milo, and T. Rauhe. Compact labeling scheme for ancestor queries. *SIAM Journal on Computing*, 35(6):1295–1309, 2006.
- [2] D. Ajwani, U. Meyer, and D. Veith. An I/O-efficient distance oracle for evolving real-world graphs. In *17th ALENEX*, pages 159–172, 2015.
- [3] T. Akiba, Y. Iwata, and Y. Yoshida. Fast exact shortest-path distance queries on large networks by pruned landmark labeling. In *32nd SIGMOD*, pages 349–360, 2013.
- [4] T. Akiba, C. Sommer, and K. Kawarabayashi. Shortest-path queries for complex networks: exploiting low tree-width outside the core. In *15th EDBT*, pages 144–155, 2012.
- [5] S. Alstrup, P. Bille, and T. Rauhe. Labeling schemes for small distances in trees. *SIAM Journal on Discrete Mathematics*, 19(2):448–462, 2005.
- [6] S. Alstrup, S. Dahlgaard, and M. B. T. Knudsen. Optimal induced universal graphs and adjacency labeling for trees. In *56th FOCS*, pages 1311–1326, 2015.
- [7] S. Alstrup, C. Gavoille, E. B. Halvorsen, and H. Petersen. Simpler, faster and shorter labels for distances in graphs. In *27th SODA*, pages 338–350, 2016.
- [8] S. Alstrup, I. L. Gørtz, E. B. Halvorsen, and E. Porat. Distance labeling schemes for trees. In *43rd ICALP*, pages 132:1–132:16, 2016.
- [9] S. Alstrup, E. B. Halvorsen, and K. G. Larsen. Near-optimal labeling schemes for nearest common ancestors. In *25th SODA*, pages 972–982, 2014.
- [10] S. Alstrup, H. Kaplan, M. Thorup, and U. Zwick. Adjacency labeling schemes and induced-universal graphs. In *47th STOC*, pages 625–634, 2015.
- [11] S. Alstrup and T. Rauhe. Small induced-universal graphs and compact implicit graph representations. In *43rd FOCS*, pages 53–62, 2002.
- [12] D. Belazzougui and G. Navarro. New lower and upper bounds for representing sequences. In *20th ESA*, pages 181–192, 2012.
- [13] N. Bonichon, C. Gavoille, and A. Labourel. Short labels by traversal and jumping. *Electronic Notes in Discrete Mathematics*, 28:153–160, 2007.
- [14] F. Chung, R. Graham, and D. Coppersmith. On trees containing all small trees. *The Theory and Applications of Graphs*, pages 265–272, 1981.
- [15] D. R. Clark. *Compact Pat Trees*. PhD thesis, University of Waterloo, 1998.
- [16] P. Elias. Universal codeword sets and representations of the integers. *IEEE transactions on information theory*, 21(2):194–203, 1975.
- [17] J. Fischer. Short labels for lowest common ancestors in trees. In *17th ESA*, pages 752–763, 2009.

- [18] P. Fraigniaud and A. Korman. Compact ancestry labeling schemes for xml trees. In *21st SODA*, pages 458–466, 2010.
- [19] C. Gavaille, M. Katz, N. A. Katz, C. Paul, and D. Peleg. Approximate distance labeling schemes. In *9th ESA*, pages 476–487, 2001.
- [20] C. Gavaille and A. Labourel. Distributed relationship schemes for trees. In *18th ISAAC*, pages 728–738, 2007. Announced at PODC’07.
- [21] C. Gavaille, D. Peleg, S. Pérennes, and R. Raz. Distance labeling in graphs. *Journal of Algorithms*, 53(1):85–112, 2004. A preliminary version in *12th SODA*, 2001.
- [22] M. Gol’dberg and E. Livshits. On minimal universal trees. *Mathematical Notes of the Academy of Sciences of the USSR*, 4(3):713–717, 1968.
- [23] G. Jacobson. Space-efficient static trees and graphs. In *30th FOCS*, pages 549–554, 1989.
- [24] S. Kannan, M. Naor, and S. Rudich. Implicit representation of graphs. *SIAM Journal on Discrete Mathematics*, 5(4):596–603, 1992.
- [25] H. Kaplan and T. Milo. Short and simple labels for small distances and other functions. In *7th WADS*, pages 246–257, 2001.
- [26] D. Peleg. Proximity-preserving labeling schemes. *Journal of Graph Theory*, 33(3):167–176, 2000.
- [27] C. Petersen, N. Rotbart, J. G. Simonsen, and C. Wulff-Nilsen. Near-optimal adjacency labeling scheme for power-law graphs. In *43rd ICALP*, pages 133:1–133:15, 2016.
- [28] M. Pătraşcu and M. Thorup. Time-space trade-offs for predecessor search. In *38th STOC*, pages 232–240, 2006.
- [29] N. G. Rotbart. *New Ideas on Labeling Schemes*. PhD thesis, University of Copenhagen, 2016.
- [30] D. D. Sleator and R. E. Tarjan. A data structure for dynamic trees. *Journal of computer and system sciences*, 26(3):362–391, 1983.
- [31] M. Thorup and U. Zwick. Compact routing schemes. In *13th SPAA*, pages 1–10, 2001.