

General Compact Labeling Schemes for Dynamic Trees

Amos Korman *

June 1, 2006

Abstract

Let F be a function on pairs of vertices. An F -labeling scheme is composed of a *marker* algorithm for labeling the vertices of a graph with short labels, coupled with a *decoder* algorithm allowing one to compute $F(u, v)$ of any two vertices u and v directly from their labels. As applications for labeling schemes concern mainly large and dynamically changing networks, it is of interest to study *distributed dynamic* labeling schemes. This paper investigates labeling schemes for dynamic trees. We consider two dynamic tree models, namely, the *leaf-dynamic* tree model in which at each step a leaf can be added to or removed from the tree and the *leaf-increasing* tree model in which the only topological event that may occur is that a leaf joins the tree.

A general method for constructing labeling schemes for dynamic trees (under the above mentioned dynamic tree models) was previously developed in [29]. This method is based on extending an existing *static* tree labeling scheme to the dynamic setting. This approach fits many natural functions on trees, such as distance, separation level, ancestry relation, routing (in both the adversary and the designer port models), nearest common ancestor etc.. Their resulting dynamic schemes incur overheads (over the static scheme) on the label size and on the communication complexity. In particular, all their schemes yield a multiplicative overhead factor of $\Omega(\log n)$ on the label sizes of the static schemes. Following [29], we develop a different general method for extending static labeling schemes to the dynamic tree settings. Our method fits the same class of tree functions. In contrast to the above paper, our trade-off is designed to minimize the label size, sometimes at the expense of communication.

Informally, for any function $k(n)$ and any static F -labeling scheme on trees, we present an F -labeling scheme on dynamic trees incurring multiplicative overhead factors (over the static scheme) of $O(\log_{k(n)} n)$ on the label size and $O(k(n) \log_{k(n)} n)$ on the amortized message complexity. In particular, by setting $k(n) = n^\epsilon$ for any $0 < \epsilon < 1$, we obtain dynamic labeling schemes with asymptotically optimal label sizes and sublinear amortized message complexity for the ancestry relation, the id-based and label-based nearest common ancestor relation and the routing function.

*Information Systems Group, Faculty of IE&M, The Technion, Haifa, 32000 Israel. E-mail: pandit@tx.technion.ac.il. Supported in part at the Technion by an Aly Kaufman fellowship.

1 Introduction

Motivation: Network representations have played an extensive and often crucial role in many domains of computer science, ranging from data structures, graph algorithms to distributed computing and communication networks. Research on network representations concerns the development of various methods and structures for cheaply storing useful information about the network and making it readily and conveniently accessible. This is particularly significant when the network is large and geographically dispersed, and information about its structure must be accessed from various local points in it. As a notable example, the basic function of a communication network, namely, message delivery, is performed by its routing scheme, which requires maintaining certain topological knowledge.

Recently, a number of studies focused on a *localized* network representation method based on assigning a (hopefully short) *label* to each vertex, allowing one to infer information about any two vertices *directly* from their labels, without using *any* additional information sources. Such labeling schemes have been developed for a variety of information types, including vertex adjacency [8, 7, 22], distance [30, 27, 19, 18, 16, 23, 35, 10, 2], tree routing [13, 36], flow and connectivity [26], tree ancestry [5, 6, 25], nearest common ancestor in trees [3] and various other tree functions, such as center, separation level, and Steiner weight of a given subset of vertices [31]. See [17] for a survey.

By now, the basic properties of localized labeling schemes for *static* (fixed topology) networks are reasonably well-understood. In most realistic contexts, however, the typical setting is highly dynamic, namely, the network topology undergoes repeated changes. Therefore, for a representation scheme to be practically useful, it should be capable of reflecting online the current up-to-date picture in a dynamic setting. Moreover, the algorithm for generating and revising the labels must be *distributed*, in contrast with the sequential and centralized label assignment algorithms described in the above cited papers.

The dynamic models investigated in this paper concern the *leaf-dynamic* tree model in which at each step a leaf can be added to or removed from the tree and the *leaf-increasing* tree model in which the only topological event that may occur is that a leaf joins the tree. We present a general method for constructing dynamic labeling schemes which is based on extending existing *static* tree labeling schemes to the dynamic setting. This approach fits a number of natural tree functions, such as routing, ancestry relation, nearest common ancestor relation, distance and separation level. Such an extension can be naively achieved by calculating the static labeling from scratch after each topological change. Though this method yields a good label size, it may incur a huge communication complexity. Another naive approach would be that each time a leaf u is added as a child of an existing node v , the label given to u is the label of v concatenated with $F(u, v)$. Such a scheme incurs very little communication, however, the labels may be huge.

Before stating the results included in this paper, we list some previous related works.

Related work: Static labeling schemes for routing on trees were investigated in [13]. For the *designer port* model, in which each node can freely enumerate its incident ports, they show how to construct a static routing scheme using labels of at most $O(\log n)$ bits on n -node trees. In the *adversary port* model, in which the port numbers are fixed by an adversary, they show how to construct a static routing scheme using labels of at most $O(\frac{\log^2 n}{\log \log n})$ bits on n -node trees. They also show that the label sizes of both schemes are asymptotically optimal. Independently, a static routing scheme for trees using $(1 + o(1)) \log n$ bit labels was introduced in [36] for the designer port model.

A static labeling scheme for the id-based nearest common ancestor (NCA) relation on trees was developed in [31] using labels of $\Theta(\log^2 n)$ bits on n -node trees. A static labeling scheme supporting the label-based NCA relation on trees using labels of $\Theta(\log n)$ bits on n -node trees is presented in [3].

In the *sequential* (non-distributed) model, dynamic data structures for trees have been studied extensively (e.g., [33, 9, 20, 4]). For comprehensive surveys on dynamic graph algorithms see [12, 15].

Labeling schemes for the ancestry relation in the leaf-dynamic tree model were investigated in [11]. They assume that once a label is given to a node it remains unchanged. Therefore, the issue of updates is not considered even for the non distributed setting. For the above model, they present a labeling scheme that uses labels of $O(m)$ bits, where m is the number of nodes added to the tree throughout the dynamic scenario. They also show that this bound is asymptotically tight. Other labeling schemes are presented in the above paper assuming that clues about the future topology of the dynamic tree are given throughout the scenario.

The study of dynamic distributed labeling schemes was initiated by [29]. Dynamic distributed distance labeling schemes on trees were investigated in [29] and [28]. In [29] they present a dynamic labeling scheme for distances in the leaf-dynamic tree model with message complexity $O(\sum_i \log^2 n_i)$, where n_i is the size of the tree when the i 'th topological event takes place. The protocol maintains $O(\log^2 n)$ bit labels, when n is the current tree size. This label size is proved in [19] to be asymptotically optimal even for the static (unweighted) trees scenario.

In [28] they develop two β -approximate distance labeling schemes (in which given two labels, one can infer a β -approximation to the distance between the corresponding nodes). The first scheme applies to the *edge-dynamic* tree model, in which the vertices of the tree are fixed but the (integer) weights of the edges may change (as long as they remain positive). The second scheme applies to the *edge-increasing* tree model, in which the only topological event that may occur is that an edge increases its weight by one. In scenarios where at most m topological events occur, the message complexities of the first and second schemes are $O(m\Lambda \log^3 n)$ and $O(m \log^3 n + n \log^2 n \log m)$, respectively, where Λ is some density parameter of the tree. The label size of both schemes is $O(\log^2 n + \log n \log W)$ where W denotes the largest edge weight in the tree.

The study of methods for extending static labeling schemes to the dynamic setting was also initiated in [29]. There, they assume the designer port model and consider two dynamic tree models, namely, the leaf-increasing and the leaf-dynamic tree models. Their approach fits a number of natural functions on trees, such as distance, separation level, ancestry relation, id-based and label-based NCA relation, routing (in both the adversary and the designer port models) etc.. Their resulting dynamic schemes incur overheads (over the static scheme) on the label size and on the communication complexity. Specifically, given a static F -labeling scheme π for trees, let $\mathcal{LS}(\pi, n)$ be the maximum number of bits in a label given by π to any vertex in any n -node tree, and let $\mathcal{MC}(\pi, n)$ be the maximum number of messages sent by π in order to assign the static labels in any n -node tree. Assuming $\mathcal{MC}(\pi, n)$ is polynomial¹ in n , the following dynamic schemes are derived. For the leaf-increasing tree model, they construct a dynamic F -labeling scheme π^{inc} . The maximum label given by π^{inc} to any vertex in any n -node tree is $O(\log n \cdot \mathcal{LS}(\pi, n))$ and the number of messages sent by π^{inc} is $O(\log n \cdot \mathcal{MC}(\pi, n))$. In the case where n_f , the final number of nodes in the tree, is known in advance, they construct a dynamic F -labeling scheme with label size $O\left(\frac{\log n_f \log n}{\log \log n_f} \cdot \mathcal{LS}(\pi, n)\right)$ and message complexity $O\left(\frac{\log n}{\log \log n_f} \cdot \mathcal{LS}(\pi, n_f)\right)$. For the leaf-dynamic tree model, they construct two dynamic F -labeling schemes. Let n_i be the size of the tree when the i 'th topological event takes place. The first dynamic F -labeling scheme has label size $O(\log n \cdot \mathcal{LS}(\pi, n))$ and message complexity $O\left(\sum_i \log n_i \cdot \frac{\mathcal{MC}(\pi, n_i)}{n_i}\right) + O(\sum_i \log^2 n_i)$ and the second dynamic F -labeling scheme has label size $O\left(\frac{\log^2 n}{\log \log n} \cdot \mathcal{LS}(\pi, n)\right)$ and message complexity $O\left(\sum_i \frac{\log n_i}{\log \log n_i} \cdot \frac{\mathcal{MC}(\pi, n_i)}{n_i}\right) + O(\sum_i \log^2 n_i)$. In particular, for all the above mentioned functions, even if n_f is known in advance, the best dynamic scheme of [29] incurs $O(\sum_i \log^2 n_i)$ message complexity and overhead of $O(\log n)$ over the label size of the corresponding static scheme.

Our contribution: Following [29], we present a different method for constructing dynamic labeling schemes in the leaf-increasing and leaf-dynamic tree model. Our method is also based on extending existing static labeling schemes to the dynamic setting. However, our resulting dynamic schemes incur different trade-offs between the overhead factors on the label sizes and the message communication. In comparison to [29], our trade-offs give better performances for the label size, sometimes at the expense of communication. Our approach fits the same class of tree functions as described in [29]. The following results apply for both the designer port model and the adversary port model. Given a static F -labeling scheme π for trees, let $\mathcal{LS}(\pi, n)$ and $\mathcal{MC}(\pi, n)$ be as before. Let $k(x)$ be any reasonable² sublinear function of x . For the leaf-increasing tree model, we construct

¹ The actual requirement is that the message complexity is bounded from above by some function f which satisfies $f(a + b) \geq f(a) + f(b)$ and $f(\Theta(n)) = \Theta(f(n))$. These two requirements are satisfied by most natural relevant functions, such as $c \cdot n^\alpha \log^\beta n$, where $c > 0$, $\alpha \geq 1$ and $\beta > 0$. For simplicity, we assume $\mathcal{MC}(\cdot, n)$ itself satisfies these requirements.

² We require that $k(x)$, $\log_{k(x)} x$ and $\frac{k(x)}{\log k(x)}$ are nondecreasing functions. Moreover we require that, $k(\Theta(x)) = \Theta(k(x))$. The above requirements are satisfied by most natural sublinear functions such as $\alpha x^\epsilon \log^\beta x$, $\alpha \log^\beta \log x$ etc..

the dynamic F -labeling scheme $\text{SDL}^{k(x)}$. The maximum number of bits in a label given by $\text{SDL}^{k(x)}$ to any vertex in any n -node tree during the dynamic scenario is $O(\log_{k(n)} n \cdot \mathcal{LS}(\pi, n))$. The maximum number of messages sent by $\text{SDL}^{k(x)}$ in any dynamic scenario is $O(k(n) \log_{k(n)} n \cdot \mathcal{MC}(\pi, n))$, where n is the final number of nodes in the tree.

In particular, by setting $k(n) = \log^\epsilon n$ for any $\epsilon > 0$, we obtain dynamic labeling schemes supporting all the above mentioned functions, with message complexity $O(n \frac{\log^{1+\epsilon} n}{\log \log n})$ and $O(\frac{\log n}{\log \log n})$ multiplicative overhead over the corresponding asymptotically optimal label size.

For the leaf-dynamic tree model, assuming $\mathcal{LS}(\pi, n)$ is multiplicative³ we construct the dynamic F -labeling scheme $\text{DL}^{k(x)}$ with the following complexities. The maximum number of bits in a label given by $\text{DL}^{k(x)}$ to any vertex in any n -node tree is $O(\log_{k(n)} n \cdot \mathcal{LS}(\pi, n))$ and the number of messages used by $\text{DL}^{k(x)}$ is $O\left(\sum_i k(n_i) (\log_{k(n_i)} n_i) \frac{\mathcal{MC}(\pi, n_i)}{n_i}\right) + O(\sum_i \log^2 n_i)$, where n_i is the size of the tree when the i 'th topological event takes place. In particular, by setting $k(n) = n^\epsilon$ for any $0 < \epsilon < 1$, we obtain dynamic labeling schemes with asymptotically the same label size as the corresponding static schemes and sublinear amortized message complexity. In particular, we get dynamic labeling schemes with sublinear amortized message complexity and asymptotically optimal label size for all the above mentioned functions. Also, by setting $k(n) = \log^\epsilon n$ for any $0 < \epsilon < 1$, we obtain dynamic labeling schemes supporting all the above mentioned functions, with message complexity $O(\sum_i \log^2 n_i)$ and $O(\frac{\log n}{\log \log n})$ multiplicative overhead over the corresponding asymptotically optimal label size. In contrast, note that for any of the above mentioned functions F , the best dynamic F -labeling scheme of [29] (in the leaf-dynamic model) has message complexity $O(\sum_i \log^2 n_i)$ and $O(\log n)$ multiplicative overhead over the corresponding asymptotically optimal label size.

Paper outline: We start with preliminaries in Section 2. In Section 3 we present the FSDL_p^k schemes which will be used in Section 4, where we introduce the dynamic labeling schemes for the leaf-increasing and the leaf-dynamic tree models. In Section 5 we discuss how to reduce the external memory used for updating and maintaining the labels.

2 Preliminaries

Our communication network model is restricted to tree topologies. The network is assumed to dynamically change via vertex additions and deletions. It is assumed that the *root* of the tree, r , is never deleted. The following types of topological events are considered.

Add-leaf: A new vertex u is *added* as a child of an existing vertex v . Subsequently, v is informed of this event.

³ We actually require that $\mathcal{LS}(\cdot, n)$ satisfies $\mathcal{LS}(\cdot, \Theta(n)) = \Theta(\mathcal{LS}(\cdot, n))$. This requirement is satisfied by most natural functions such as $c \cdot n^\alpha \log^\beta n$, where $c > 0$, $\alpha \geq 0$ and $\beta > 0$.

Remove-leaf: A leaf of the tree is *deleted*. Subsequently, the leaf’s parent is informed of this event.

We consider two types of dynamic models. Namely, the *leaf-increasing* tree model in which the only topological event that may occur is of type add-leaf, and the *leaf-dynamic* tree model in which both types of topological events may occur.

Incoming and outgoing links at every node are identified by so called *port-numbers*. When a new child is added to a node v , the corresponding ports are assigned a unique port-number, in the sense that no currently existing two ports of v have the same port-number. We consider two main variations, namely, the *designer port* model and the *adversary port* model. The former allows each node v to freely enumerate its incident ports while the latter assumes that the port numbers are fixed by an adversary.

Our method is applicable to any function F such that for every two vertices u and v in the tree the following condition is satisfied.

- (C1) For every vertex w on the path between u and v , $F(u, v)$ can be calculated in polynomial time from $F(u, w)$ and $F(w, v)$.

In particular, our method can be applied to the ancestry relation, the id-based and label-based NCA relations and for the distance, separation level and routing functions (both in the designer and the adversary port models), thereby extending static labeling schemes such as those of [3, 13, 36, 31, 30] to the dynamic setting. We further assume, for simplicity of presentation, that F is symmetric, i.e., $F(u, v) = F(v, u)$. A slight change to the suggested protocols handles the more general case, without affecting the asymptotic complexity results.

A *labeling scheme* $\pi = \langle \mathcal{M}_\pi, \mathcal{D}_\pi \rangle$ for a function F on pairs of vertices of a tree is composed of the following components:

1. A *marker* algorithm \mathcal{M}_π that given a tree, assigns labels to its vertices.
2. A polynomial time *decoder* algorithm \mathcal{D}_π that given the labels $L(u)$ and $L(v)$ of two vertices u and v , outputs $F(u, v)$.

In this paper we are interested in distributed networks where each vertex in the tree is a processor. This does not affect the definition of the decoder algorithm of the labeling scheme since it is performed locally, but the marker algorithm changes into a *distributed marker protocol*.

Let us first consider static networks, where no changes in the topology of the network are allowed. For these networks we define *static* labeling schemes, where the marker protocol \mathcal{M} is initiated at the root of a tree network and assigns static labels to all the vertices once and for all.

We use the following complexity measures to evaluate a static labeling scheme $\pi = \langle \mathcal{M}_\pi, \mathcal{D}_\pi \rangle$.

1. *Label Size*, $\mathcal{LS}(\mathcal{M}_\pi, n)$: the maximum number of bits in a label assigned by \mathcal{M}_π to any vertex on any n -vertex tree.
2. *Message Complexity*, $\mathcal{MC}(\mathcal{M}_\pi, n)$: the maximum number of messages sent by \mathcal{M}_π during the labeling process on any n -vertex tree. (Note that messages can only be sent between neighboring vertices).

We assume that the static labeling scheme assigns unique labels. For any static labeling scheme, this additional requirement can be ensured at an extra additive cost of at most n to $\mathcal{MC}(n)$ and $\log n$ to $\mathcal{LS}(n)$.

Example 2.1 *The following is a possible static labeling scheme STATDFS for the ancestry relation on trees based on the notion of interval schemes ([34], cf. [32]). Given a rooted tree, simply perform a depth-first search starting at the root, assigning each vertex v the interval $I(v) = [a, b]$ where a is its DFS number and b is the largest DFS number given to any of its descendants. The corresponding decoder decides that v is an ancestor of w iff their corresponding intervals, $I(v)$ and $I(w)$, satisfy $I(v) \subseteq I(w)$. It is easy to verify that this is a correct labeling scheme for the ancestry relation. Clearly, $\mathcal{MC}(\text{STATDFS}, n) = O(n)$ and $\mathcal{LS}(\text{STATDFS}, n) = O(\log n)$.*

Labeling schemes for routing are presented in [13]. They consider both the designer port model and the adversary port model. The schemes of [13] are designed as a sequential algorithm, but examining the details reveals that these algorithms can be easily transformed into distributed protocols. In the designer port model, we get a static labeling scheme for routing with label size and message complexity similar to those of the STATDFS static labeling scheme. In the adversary port model we get a static labeling scheme for routing with linear communication and $O(\frac{\log^2 n}{\log \log n})$ label size. The label sizes of both schemes are asymptotically optimal.

The dynamic labeling schemes involve a marker protocol \mathcal{M} which is activated after every change in the network topology. The protocol \mathcal{M} maintains the labels of all vertices in the underlying graph so that the corresponding decoder algorithm will work correctly. We assume that the topological changes occur serially and are sufficiently spaced so that the protocol has enough time to complete its operation in response to a given topological change before the occurrence of the next change.

We distinguish between the label $\mathcal{M}(v)$ given to each node v to deduce the required information in response to online queries, and the additional external storage $Memory(v)$ at each node v , used during updates and maintenance operations. For certain applications (and particularly routing), the label $\mathcal{M}(v)$ is often kept in the router itself, whereas the additional storage $Memory(v)$ is kept on some external storage device. Subsequently, the size of $\mathcal{M}(v)$ seems to be a more critical consideration than the total amount of storage needed for the information maintenance.

For the leaf-increasing tree model, we use the following complexity measures to evaluate a dynamic labeling scheme $\pi = \langle \mathcal{M}_\pi, \mathcal{D}_\pi \rangle$.

1. *Label Size*, $\mathcal{LS}(\mathcal{M}_\pi, n)$: the maximum size of a label assigned by the marker protocol \mathcal{M}_π to any vertex on any n -vertex tree in any dynamic scenario.
2. *Message Complexity*, $\mathcal{MC}(\mathcal{M}_\pi, n)$: the maximum number of messages sent by \mathcal{M}_π during the labeling process in any scenario where n is the final number of vertices in the tree.

Finally, we consider the leaf-dynamic tree model, where both additions and deletions of vertices are allowed. Instead of measuring the message complexity in terms of the maximal number of nodes in the scenario, for more explicit time references, we use the notation $\bar{n} = (n_1, n_2, \dots, n_f)$ where n_i is the size of the tree immediately after the i 'th topological event takes place. For simplicity, we assume $n_1 = 1$ unless stated otherwise. The definition of $\mathcal{LS}(\mathcal{M}_\pi, n)$ remains as before, and the definition of the message complexity changes into the following.

Message Complexity, $\mathcal{MC}(\mathcal{M}_\pi, \bar{n})$: the maximum number of messages sent by \mathcal{M}_π during the labeling process in any scenario where n_i is the size of the tree immediately after the i 'th topological event takes place.

3 The finite semi-dynamic F -labeling schemes FSDL_p^k

In this section, we consider the leaf-increasing tree model and assume that the initial tree contains a single vertex, namely, its root. Given a static F -labeling scheme $\pi = \langle \mathcal{M}_\pi, \mathcal{D}_\pi \rangle$, we first fix some integer k and then, for each integer $p \geq 1$, we recursively define the dynamic scheme FSDL_p^k which acts on growing trees and terminates at some point. Each dynamic scheme FSDL_p^k is guaranteed to function as a dynamic F -labeling scheme as long as it operates. It will follow that Scheme FSDL_p^k terminates only when n , the number of nodes in the current tree, is at least k^p . Moreover, the overheads (over π) of Scheme FSDL_p^k are $O(p)$ on the label size and $O(p \cdot k)$ on the message complexity. The FSDL_p^k schemes are used in the next section as building blocks for our dynamic F -labeling schemes. Let us first give an informal description of the FSDL_p^k scheme and its analysis.

3.1 Overview of Scheme FSDL_p^k

Scheme FSDL_p^k repeatedly invokes a *reset* operation on different subtrees, in which the marker protocol of the static labeling scheme is applied and the labels it produces are used to construct the dynamic labels. It will follow that just before Scheme FSDL_p^k terminates, a reset operation is invoked on the whole current tree.

The FSDL_p^k schemes are defined recursively on p as follows. In Scheme FSDL_1^k , whenever a new vertex joins the tree, a reset operation is invoked on the whole tree, in which each vertex receives the label given to it by the marker protocol of the static labeling scheme. The decoder of Scheme FSDL_1^k is simply the decoder algorithm of static labeling scheme. Using a counter at the root, after k such reset operations, the scheme terminates.

Given Scheme FSDL_p^k , we now define Scheme FSDL_{p+1}^k . We start by running Scheme FSDL_p^k at the root, until it is supposed to terminate. As mentioned before, just before Scheme FSDL_p^k terminates, a reset operation is invoked on T_0 , the whole current tree. This reset operation is referred to as a $(p+1)$ -global reset operation (it may also be referred to as an l -global reset operation for other l 's). Before this $(p+1)$ -global reset operation, the FSDL_{p+1}^k scheme is simply the FSDL_p^k scheme (which is applied at the root). I.e., the label given to any vertex v by the FSDL_{p+1}^k scheme is the label given to v by the FSDL_p^k scheme, and the decoder of Scheme FSDL_{p+1}^k is simply the decoder of Scheme FSDL_p^k . During the above mentioned $(p+1)$ -global reset operation, each vertex $v \in T_0$ receives the label $\mathcal{M}_\pi(v)$ given to v by the marker algorithm of the static labeling scheme. Instead of terminating Scheme FSDL_p^k , we continue as follows. For every $v \in T_0$, let T_v denote the dynamic subtree rooted at v that contains v and v 's future children as well as all their future descendants. After the above mentioned $(p+1)$ -global reset operation, each vertex $v \in T_0$ invokes Scheme FSDL_p^k on T_v . If, at some point, one of these FSDL_p^k schemes is supposed to terminate, instead of terminating it, a reset operation (which is also referred to as a $(p+1)$ -global reset operation) is invoked on T_0 , the whole current tree. Again, after the above mentioned $(p+1)$ -global reset operation, each vertex $v \in T_0$ invokes Scheme FSDL_p^k on T_v . As before, if, at some point, one of these FSDL_p^k schemes is supposed to terminate, instead of terminating it, a $((p+1)$ -global) reset operation is invoked on the whole current tree, and so forth. Using a counter at the root, after k such $(p+1)$ -global reset operations, the FSDL_{p+1}^k scheme terminates.

After any of the above mentioned $(p+1)$ -global reset operations, the label given to a vertex $w \in T_v$ for some $v \in T_0$ contains the following components. The label $\mathcal{M}_\pi(v)$, the relation $F(w, v)$ and the label given to w by the FSDL_p^k scheme that is applied on T_v . Given the labels $L(x)$ and $L(y)$ of two vertices $x \in T_v$ and $y \in T_u$, where $v \neq u$, the decoder algorithm finds $F(x, y)$ using **1**) the static decoder algorithm applied on $\mathcal{M}_\pi(v)$ and $\mathcal{M}_\pi(u)$, **2**) the relations $F(x, v)$ and $F(y, u)$ and **3**) the condition C1. If x and y are at the same subtree T_v , then the decoder finds $F(x, y)$ using the decoder algorithm of the FSDL_p^k scheme applied on the labels given to x and y by the FSDL_p^k scheme (which was invoked on T_v).

Using induction on p , it follows that Scheme FSDL_p^k may terminate only when the number of nodes in the tree is at least k^p . Also, using induction on p , it can be shown that the label size of the dynamic scheme is at most $O(p)$ times the label size of the static scheme π . The fact that the message complexity of Scheme FSDL_p^k is $O(p \cdot k)$ times the message complexity of π , intuitively follows from the following facts. 1) for every $1 \leq l \leq p$, the different applications of Scheme FSDL_l^k act on edge disjoint subtrees and 2) for every $1 \leq l \leq p$, every application of Scheme FSDL_l^k invokes an l -global reset operations at most k times.

Scheme FSDL_p^k invokes Scheme FSDL_l^k for different l 's on different subtrees. These different applications of Scheme FSDL_l^k induce a decomposition of the tree into subtrees of different levels; an l -level subtree is a subtree on which Scheme FSDL_l^k is invoked. In particular, the whole tree is a p -level subtree and each vertex is contained in precisely one l -level subtree, for each

$1 \leq l \leq p$. Moreover, subtrees of the same level are edge-disjoint, however, subtrees of different levels may overlap, in particular, for $1 \leq l < p$, any l -level subtree is (not necessarily strictly) contained in some $l + 1$ -level subtree. Note that l -global reset operations can be applied only on l -level subtrees. The above mentioned decomposition of the tree into subtrees is referred to as *the subtrees decomposition*. As shown later, the subtrees decomposition is quite different from the tree decomposition (into bubbles) of [29], on which their dynamic schemes are based upon.

In order to add intuition, we now give a short informal description of the FSDL_p^k scheme and the subtrees decomposition from a non-recursive point of view. Initially, the root is considered as an l -level subtree for every $1 \leq l \leq p$. At any time, given the current subtrees decomposition, Scheme FSDL_p^k operates as follows. Whenever a leaf v joins the tree as a child of vertex u , for every $1 \leq l \leq p$, the l -level subtree $T_l(u)$ becomes $T_l(u) \cup \{v\}$ and $T_l(v)$ is defined to be $T_l(u)$. In addition, a (1-global) reset operation is invoked on $T_1(v)$. This 1-global reset operation may result in a sequence of reset operations as follows. If, after the last reset operation, the root of $T_1(v)$ went through k (1-global) reset operations then the following happen.

- 1) If, just before the reset operation, $T_2(v)$ strictly contained $T_1(v)$, then a (2-global) reset operation is invoked on $T_2(v)$,
- 2) $T_2(v)$ remains a 2-level subtree and $T_1(v)$ is no longer considered as part of the subtrees decomposition,
- 3) each vertex $w \in T_2(v)$ becomes the root of a new 1-level subtree, namely T_w .

In general, for every $1 \leq l \leq p - 1$, if after the last l -global reset operation, the root of $T_l(v)$ went through k (l -global) reset operations then the following happen.

- 1) If, just before the reset operation, $T_{l+1}(v)$ strictly contained $T_l(v)$, then an $((l + 1)$ -global) reset operation is invoked on $T_{l+1}(v)$,
- 2) $T_{l+1}(v)$ remains an $(l + 1)$ -level subtree but for every $1 \leq l' \leq l$, every subtree in the subtrees decomposition containing an edge of $T_{l+1}(v)$ is removed from the subtrees decomposition,
- 3) for every $1 \leq l' \leq l$ and every vertex $w \in T_{l+1}(v)$, the subtree T_w is added to the subtrees decomposition as a new l' -level subtree.

If, after the last p -global reset operation, T went through k (p -global) reset operation then Scheme FSDL_p^k terminates.

We are now ready to describe the FSDL_p^k scheme more formally.

3.2 Scheme FSDL_p^k

We start with the following definition. A *finite semi-dynamic F -labeling scheme* is a dynamic F -labeling scheme that is applied on a dynamically growing tree T and terminates at some point. I.e.,

the root can be in one of two states, namely, 0 or 1, where initially, the root is in state 1 and when the root changes its state to 0, the scheme is considered to be terminated. The requirement from a finite semi-dynamic F -labeling scheme is that until the root changes its state to 0, the scheme operates as a dynamic F -labeling scheme. For a finite semi-dynamic F -labeling scheme, \mathcal{S} , we define its *stopping time* $ST(\mathcal{S})$ to be the minimum number of nodes that have joined the tree until the time \mathcal{S} terminates, taken over all scenarios. Assuming $ST(\mathcal{S}) \geq n$, the complexities $\mathcal{LS}(\mathcal{S}, n)$ and $\mathcal{MC}(\mathcal{S}, n)$ are defined in the same manner as they are defined for dynamic labeling schemes.

Let $\pi = \langle \mathcal{M}_\pi \mathcal{D}_\pi \rangle$ be a static F -labeling scheme such that $\mathcal{MC}(\pi, n)$ is polynomial in n (see footnote 1). Fix some integer $k > 1$. We now describe for each integer $p \geq 1$, the finite semi-dynamic F -labeling scheme $\text{FSDL}_p^k = \langle \mathcal{M}_p, \mathcal{D}_p \rangle$.

Our dynamic schemes repeatedly engage the marker protocol of the static labeling scheme, and use the labels it produces to construct the dynamic labels. In doing so, the schemes occasionally apply to the already labeled portion of the tree a *reset* operation (defined below) invoked on some subtree T' .

Sub-protocol $\text{RESET}(T')$

- The root of T' initiates broadcast and convergcast operations (see [32]) in order to calculate $n(T')$, the number of vertices in T' .
- The root of T' invokes the static labeling scheme π on T' .

We describe the finite semi-dynamic F -labeling schemes FSDL_p^k in a recursive manner. It will follow from our description that Scheme FSDL_p^k terminates immediately after some Sub-protocol RESET is invoked on the whole current tree, T . Throughout the run of Scheme FSDL_p^k , the root r keeps a counter μ_p . We start by describing FSDL_1^k .

Scheme FSDL_1^k

1. If a new node joins as a child of the root r then r invokes Sub-protocol $\text{RESET}(T)$ on the current tree (which contains two vertices).
2. The root initializes its counter to $\mu_1 = 1$.
3. If a new node joins the tree, it sends a signal to r instructing it to invoke Sub-protocol $\text{RESET}(T)$ on the current tree T .
4. The root r sets $\mu_1 = \mu_1 + 1$. If $\mu_1 = k$ then r changes its state to 0 and the scheme terminates. Otherwise we proceed by going back to the previous step.

Clearly FSDL_1^k is a finite semi-dynamic F -labeling scheme.

Given the finite semi-dynamic F -labeling scheme $\text{FSDL}_{p-1}^k = \langle \mathcal{M}_{p-1}, \mathcal{D}_{p-1} \rangle$, we now describe the scheme $\text{FSDL}_p^k = \langle \mathcal{M}_p, \mathcal{D}_p \rangle$.

Scheme FSDL_p^k

1. We first initiate FSDL_{p-1}^k at r . At some point during the scenario, (after some application of Sub-protocol $\text{RESET}(T)$), the root is supposed to change its state to 0 in order to terminate Scheme FSDL_{p-1}^k . Instead of doing so, we proceed to Step 2.
2. The root initializes its counter to $\mu_p = 1$.
3. Let T_0 be the tree at the last time Sub-protocol RESET was applied and let $\mathcal{M}_\pi(u)$ be the static label given to $u \in T_0$ in the second step of that sub-protocol.
4. If $\mu_p = k$ then the root changes its state to 0 and the scheme terminates. Otherwise we continue to the next step.
5. The root broadcasts a signal to all the vertices in T_0 instructing each vertex u to invoke Scheme FSDL_{p-1}^k on T_u , the future subtree rooted at u which contains u and u 's future children as well as their future descendants. Let $\text{FSDL}_{p-1}(u) = \langle \mathcal{M}_{p-1}^u, \mathcal{D}_{p-1} \rangle$ denote the scheme FSDL_{p-1}^k which is invoked by u .
6. For each vertex w , let u be the vertex in T_0 such that $w \in T_u$. The label given to w by the marker \mathcal{M}_p is defined as $\mathcal{M}_p(w) = \langle \mathcal{M}_\pi(u), F(u, w), \mathcal{M}_{p-1}^u(w) \rangle$.
7. For a vertex z and $i \in \{1, 2, 3\}$, let $L_i(z)$ denote the i 'th field of $L(z)$. Given two labels $L(x)$ and $L(y)$ of two vertices x and y , the decoder \mathcal{D}_p operates as follows.
 - If $L_1(x) = L_1(y)$ (which means that x and y belong to the same subtree T_u for some $u \in T_0$) then \mathcal{D}_p outputs $\mathcal{D}_{p-1}(L_3(x), L_3(y))$.
 - If $L_1(x) \neq L_1(y)$ then this means that $x \in T_u$ and $y \in T_v$ where both u and v belong to T_0 . Furthermore, u is on the path from x to v and v is on the path from x to y . Therefore $F(x, u) = L_2(x)$, $F(u, v) = \mathcal{D}_\pi(L_1(x), L_1(y))$ and $F(v, y) = L_2(y)$. The decoder proceeds using Condition (C1) on F .
8. If at some point during the scenario, some vertex $u \in T_0$ is supposed to terminate $\text{FSDL}_{p-1}(u)$ by changing its state to 0, then instead of doing so, it sends a signal to the root r which in turn invokes Sub-protocol $\text{RESET}(T)$ and sets $\mu_p = \mu_p + 1$. We proceed by going back to Step 3.

By induction it is easy to show that Scheme FSDL_p^k is indeed a finite semi-dynamic F -labeling scheme. Let us first prove that the stopping time of Scheme FSDL_p^k is at least k^p .

Claim 3.1 $\mathcal{ST}(\text{FSDL}_p^k) \geq k^p$.

Proof: We prove the claim by induction on p . For $p = 1$, it is clear from the description of Scheme FSDL_1^k that if this scheme terminates then the number of nodes that have joined the tree is k . Assume by induction that $\mathcal{ST}(\text{FSDL}_p^k) \geq k^p$ and consider Scheme FSDL_{p+1}^k .

Recall that Scheme FSDL_{p+1}^k initially invokes (in Step 1) Scheme FSDL_p^k until the latter is supposed to terminate. Then (after some messages are sent), by Step 5 of Scheme FSDL_{p+1}^k , each vertex on the current tree invokes Scheme FSDL_p^k on its future subtree until one of these schemes is supposed to terminate. If at this point, Scheme FSDL_{p+1}^k does not terminate, then again (after some messages are sent), each vertex on the current tree invokes Scheme FSDL_p^k on its future subtrees, on so forth. By Steps 2,4 and 8 of Scheme FSDL_{p+1}^k , when Scheme FSDL_{p+1}^k terminates, Step 5 has been applied $k - 1$ times and Step 1 has been applied once. In each of these applications of Scheme FSDL_p^k (which act on disjoint sets of edges), by our induction hypothesis, at least k^p vertices have joined the corresponding subtree. Altogether, we obtain that at least k^{p+1} vertices have joined the tree. The claim follows. \blacksquare

Lemma 3.2 $\bullet \mathcal{LS}(\text{FSDL}_p^k, n) = O(p \cdot \mathcal{LS}(\pi, n))$.

$\bullet \mathcal{MC}(\text{FSDL}_p^k, n) \leq 5pk \cdot \mathcal{MC}(\pi, n)$.

Proof: The existence of a static F -labeling scheme π with labels of at most $\mathcal{LS}(\pi, n)$ bits implies that for any two vertices u and v in any n -node tree, $F(u, v)$ can be encoded using $O(\mathcal{LS}(\pi, n))$ bits. This can be done by simply writing the labels of the two vertices. The first part of the lemma follows by induction. We now turn to prove the second part of the lemma using induction on p .

Using the fact that $\mathcal{MC}(\pi, a) \geq a$ for every $a \geq 1$, it follows that for $p = 1$, $\mathcal{MC}(\text{FSDL}_1^k) \leq 5k \cdot \mathcal{MC}(\pi, n)$. Assume by induction that $\mathcal{MC}(\text{FSDL}_p^k) \leq 5pk \cdot \mathcal{MC}(\pi, n)$ and consider Scheme FSDL_{p+1}^k . We distinguish between two types of messages sent by Scheme FSDL_{p+1}^k during the dynamic scenario. The first type of messages consists of the messages sent in the different applications of Scheme FSDL_p^k . The second type of messages consists of the broadcast messages in Step 5 of Scheme FSDL_{p+1}^k and the messages resulted from the applications of Step 8 of Scheme FSDL_{p+1}^k (which correspond to sending a signal to the root and applying Sub-protocol RESET). Let us first bound from above the number of messages of the first type. Recall that Scheme FSDL_{p+1}^k initially invokes Scheme FSDL_p^k until the latter is supposed to terminate. Then messages of the second type are sent and then each vertex on the current tree invokes Scheme FSDL_p^k on its future subtree until one of these schemes is supposed to terminate. Again, if at this point, Scheme FSDL_{p+1}^k does not terminate, then messages of the second type are sent and then each vertex on the current tree invokes Scheme FSDL_p^k on its future subtrees, on so forth. Note that the different applications of FSDL_p^k act on disjoint sets of edges and since we assume that $\mathcal{MC}(\pi, (a+b)) \geq \mathcal{MC}(\pi, a) + \mathcal{MC}(\pi, b)$ is satisfied for every $a, b \geq 1$, we obtain (by our induction hypothesis) that the number of messages of the first type is at most $5pk \cdot \mathcal{MC}(\pi, n)$.

By Steps 2,4 and 8 of Scheme FSDL_{p+1}^k we get that Step 3 of Scheme FSDL_{p+1}^k can be applied

at most k times. Using the fact that $\mathcal{MC}(\pi, a) \geq a$ for every $a \geq 1$, the total number of messages of the second type sent by Scheme FSDL_{p+1}^k is at most $5k\mathcal{MC}(\pi, n)$. Altogether, we obtain that the number of messages sent by Scheme FSDL_{p+1}^k during the dynamic scenario is at most $5pk \cdot \mathcal{MC}(\pi, n) + 5k \cdot \mathcal{MC}(\pi, n) = 5(p+1)k \cdot \mathcal{MC}(\pi, n)$. The second part of the lemma follows. \blacksquare

3.3 The subtrees decomposition

We refer to the a reset operations mentioned in either Step 1 or Step 8 of the description of Scheme FSDL_p^k as a *p-global reset operation*. Scheme FSDL_p^k invokes Scheme FSDL_l^k for different l 's (where $1 \leq l \leq p-1$) on different subtrees. These different applications of Scheme FSDL_l^k induce a decomposition of the tree into subtrees of different levels as follows. At any time during the dynamic scenario, the whole tree is considered as a p -level subtree. At any time before the first p -global reset operation, the whole tree is also considered as a $(p-1)$ -level subtree. At any given time after the first p -global reset operation, let T_0 denote the tree during the last p -global reset operation. Between any two p -global reset operations, the edges of T_0 are not considered as part of any l -level subtree, where $l < p$, in other words, the edges of T_0 are only considered as part of the p -level subtree, which is the whole tree. However, for each $v \in T_0$, the dynamic subtree T_v is now considered as a $(p-1)$ -level subtree. The decomposition into subtrees induced by Scheme FSDL_p^k continues recursively using the decomposition into subtrees induced by the FSDL_{p-1}^k schemes which are applied on T_v for every $v \in T_0$. We refer to the resulting decomposition as *the subtrees decomposition*. The following properties easily follow from the description of Scheme FSDL_p^k .

Subtrees decomposition properties

1. For any given $1 \leq l \leq p$, the l -level subtrees are edge disjoint.
2. For every $1 \leq l \leq p$, each vertex v belongs to precisely one l -level subtree; we denote this subtree by $T_l(v)$.
3. Subtrees of different levels may overlap, in particular, for $1 \leq l < p$, any l -level subtree is (not necessarily strictly) contained in some $l+1$ -level subtree.
4. If v is not the root of $T_l(v)$, then all v 's descendants also belong to $T_l(v)$.
5. Each reset operation may only be invoked on subtrees of the subtrees decomposition.

We note that, the dynamic schemes of [29] are based on the *bubble tree decomposition* (see Subsection 4.1.1 of [29]). As can be observed by the above properties, the subtrees decomposition is quite different from the bubble tree decomposition of [29].

Since reset operations are carried on the subtrees of the subtrees decomposition, each vertex v must ‘know’, for each $1 \leq l \leq p$, which of its incident edges belong to $T_l(v)$. The method by which each vertex v implements the above is discussed in Section 5.

4 The dynamic F -labeling schemes

Let $\pi = \langle \mathcal{M}_\pi \mathcal{D}_\pi \rangle$ be a static F -labeling scheme such that $\mathcal{MC}(\pi, n)$ is polynomial in n (see footnote 1) and let $k(x)$ be a sublinear function (see footnote 2). We first construct the dynamic F -labeling scheme $\text{SDL}^{k(x)}$ for the leaf-increasing tree model and then show how to transform it to our dynamic F -labeling scheme $\text{DL}^{k(x)}$ which is applicable in the leaf-dynamic tree model.

4.1 The dynamic F -labeling scheme $\text{SDL}^{k(x)}$

We now describe our dynamic F -labeling scheme $\text{SDL}^{k(x)}$ which operates in the leaf-increasing tree model. Scheme $\text{SDL}^{k(x)}$ invokes the FSDL_p^k schemes for different parameters k and p . Let us first describe the case in which the initial tree contains a single vertex, i.e., its root. In this case, Scheme $\text{SDL}^{k(x)}$ operates as follows.

Scheme $\text{SDL}^{k(x)}$

1. Invoke Scheme $\text{FSDL}_1^{k(1)}$.
2. Recall that while invoking Scheme FSDL_p^k , just before this scheme is supposed to terminate, Sub-protocol $\text{RESET}(T)$ is invoked in which n' , the number of nodes in T , is calculated. For such n' , let p' be such that $k(n')^{p'} \leq 2 \cdot n' < k(n')^{p'+1}$. Let $p = p' + 2$ and let $k = k(n')$. Instead of terminating the above scheme, we proceed to the next step, i.e., Step 3 in Scheme $\text{SDL}^{k(x)}$.
3. The root of the whole tree invokes Scheme FSDL_p^k (with the parameters k and p defined in the previous step) while ignoring Step 1 of that scheme, i.e., start directly in Step 2 of Scheme FSDL_p^k .
At some point, Scheme FSDL_p^k is supposed to terminate. Instead of terminating it, we proceed by going back to Step 2 of Scheme $\text{SDL}^{k(x)}$.

Theorem 4.1 $\text{SDL}^{k(x)}$ is a dynamic F -labeling scheme for the leaf-increasing tree model with the following complexities.

- $\mathcal{LS}(\text{SDL}^{k(x)}, n) = O(\log_{k(n)} n \cdot \mathcal{LS}(\pi, n))$.
- $\mathcal{MC}(\text{SDL}^{k(x)}, n) = O(k(n)(\log_{k(n)} n)\mathcal{MC}(\pi, n))$.

Proof: At any given time t , there exist constants k and p such that Scheme FSDL_p^k is applied by Scheme $\text{SDL}^{k(x)}$. Let n be the current number of nodes in the tree and let n' be the number of nodes in the tree when the current Scheme FSDL_p^k was initiated. We have $k^{p-2}(n') \leq 2 \cdot n' < k^{p-1}(n')$ and therefore $p-2 \leq \log_{k(n')} 2n'$. Since $n' \leq n$ then by assumptions on $k(x)$, we get that $p = O(\log_{k(n)} n)$ and the first part of the theorem follows from the first part of Lemma 3.2. We now turn to prove the second part of the theorem.

For analysis purposes, we divide the scenario into sub-scenarios according to the different applications of Step 3 in Scheme $\text{SDL}^{k(x)}$. We define these sub-scenarios as follows. Recall that initially, Scheme $\text{SDL}^{k(x)}$ invokes Scheme $\text{FSDL}_1^{k(1)}$ until the latter is supposed to terminate. We refer to the above mentioned scenario as the 1'st scenario. For $i > 1$, the i 'th scenario corresponds to the scenario between the $i-1$ 'st and the i 'th applications of Step 3 of Scheme $\text{SDL}^{k(x)}$ (the i 'th scenario includes the $i-1$ 'st application of Step 3 and does not include the i 'th application of Step 3). Let k_i and p_i be the parameters of the FSDL scheme corresponding to the i 'th scenario and denote this scheme by Scheme $\text{FSDL}_{p_i}^{k_i}$. Let n_i be the number of nodes in the tree at the beginning of the i 'th scenario.

Claim 1: For every $i > 1$, $p_{i-1} \cdot k_{i-1} \leq (p_i - 1) \cdot k_i$.

Proof: Since Scheme $\text{FSDL}_{p_{i-1}}^{k_{i-1}}$ was supposed to terminate when Scheme $\text{FSDL}_{p_i}^{k_i}$ was initiated, Step 5 in Scheme $\text{FSDL}_{p_{i-1}}^{k_{i-1}}$ was applied $k_{i-1} - 1$ times. Therefore, by Claim 3.1, we obtain $(k_{i-1} - 1) \cdot k_{i-1}^{p_{i-1}-1} \leq n_i$ and therefore $k_{i-1}^{p_{i-1}} \leq 2 \cdot n_i$, which implies $p_{i-1} \leq \log_{k_{i-1}} 2n_i$. We therefore get that $p_{i-1} \cdot k_{i-1} \leq k_{i-1} \log_{k_{i-1}} 2n_i = \frac{k_{i-1}}{\log k_{i-1}} \log 2n_i$. By our assumption on $k(x)$, we get that $\frac{k_{i-1}}{\log k_{i-1}} \log 2n_i \leq \frac{k_i}{\log k_i} \log 2n_i = k_i \log_{k_i} 2n_i$. By the choice of k_i and p_i , We obtain that $2n_i < k_i^{p_i-1}$ and therefore $\log_{k_i} 2n_i < p_i - 1$. Altogether, we obtain $p_{i-1} \cdot k_{i-1} \leq (p_i - 1) \cdot k_i$, as desired. ■

Claim 2: For any i , at any given time t during the i 'th scenario, if the number of nodes in the tree at time t is n , then the total number of messages sent by $\text{SDL}^{k(x)}$ until time t is at most $5k_i p_i \mathcal{MC}(\pi, n)$.

Proof: We prove the claim by induction on i . For $i = 1$ we have $k_1 = k(1)$ and $p_1 = 1$ and the claim follows by the second part of Lemma 3.2. Assume that the claim is true for $i-1$ and consider a time t in the i 'th scenario such that the number of nodes in the tree at time t is n .

We distinguish between three types of messages sent until time t . The first type of messages consists of the messages sent until the i 'th scenario was initiated. The second type of messages consists of the messages sent in the different applications of Scheme $\text{FSDL}_{p_{i-1}}^{k_{i-1}}$ in Step 5 of Scheme $\text{FSDL}_{p_i}^{k_i}$. The third type of messages consists of the broadcast messages resulting from Step 5 of Scheme $\text{FSDL}_{p_i}^{k_i}$ and the messages sent during the applications of Step 8 of Scheme $\text{FSDL}_{p_i}^{k_i}$ (which correspond to sending a signal to the root and applying Sub-protocol RESET).

By our induction hypothesis and the previous claim, the number of messages of the first type is at most $5k_{i-1} p_{i-1} \mathcal{MC}(\pi, n_i) \leq 5k_i (p_i - 1) \mathcal{MC}(\pi, n_i)$.

By the second part of Lemma 3.2, we get that if $\text{FSDL}_{p_i-1}^{k_i}$ is invoked on a growing tree whose current number of nodes is n' , then the number of messages sent by $\text{FSDL}_{p_i-1}^{k_i}$ is at most $5k_i(p_i - 1)\mathcal{MC}(\pi, n')$. Using similar arguments as in the proof of Lemma 3.2, by our assumptions on $\mathcal{MC}(\pi, \cdot)$, we obtain that the total number of messages of both the first type and the second type is at most $5k_i(p_i - 1)\mathcal{MC}(\pi, n)$. Moreover, since Step 3 of Scheme $\text{FSDL}_{p_i}^{k_i}$ is applied at most k_i times and since $\mathcal{MC}(\pi, a) \geq a$ for every $a \geq 1$, the total number of messages of the third type is at most $5k_i\mathcal{MC}(\pi, n)$. Altogether, we get that the number of messages sent by time t is at most $5k_i(p_i - 1)\mathcal{MC}(\pi, n) + 5k_i\mathcal{MC}(\pi, n) = 5k_i p_i \mathcal{MC}(\pi, n)$ and the claim follows. ■

Fix a time t and let n be the number of nodes in the tree at time t . Let i be such that time t belongs to the i 'th scenario. By the choice of k_i and p_i and by our assumptions on $k(x)$, we have $k_i = k(n_i) \leq k(n)$ and $p_i = O(\log_{k(n_i)} n_i) = O(\log_{k(n)} n)$. The second part of the theorem follows from Claim 2. ■

Let us now describe how to extend $\text{SDL}^{k(x)}$ to the scenario in which the initial tree T_0 does not necessarily contains just the root. In this case, Scheme $\text{SDL}^{k(x)}$ operates as follows.

Scheme $\text{SDL}^{k(x)}$ initiated on T_0

1. The root of T_0 invokes Sub-protocol $\text{RESET}(T_0)$ in which the number of nodes n_0 in the initial tree is calculated. Let p' be such that $k(n_0)^{p'} \leq 2 \cdot n_0 < k(n_0)^{p'+1}$. Let $k = k(n_0)$ and let $p = p' + 2$.
2. The root invokes Scheme FSDL_p^k while ignoring Step 1 of that scheme, i.e., start directly in Step 2 of Scheme FSDL_p^k . At some point, Scheme FSDL_p^k is supposed to terminate. Instead of terminating it, we proceed by going to the next step, i.e., Step 3 of this scheme.
3. Recall that while invoking Scheme FSDL_p^k , just before this scheme is supposed to terminate, Sub-protocol $\text{RESET}(T)$ is invoked in which n' , the number of nodes in T , is calculated. For such n' , let p' be such that $k(n')^{p'} \leq 2 \cdot n' < k(n')^{p'+1}$. Let $p = p' + 2$ and let $k = k(n')$. Instead of terminating the above scheme, we proceed by going back to the previous step, i.e., Step 2.

The proof of the following theorem follows similar steps as the proof of Theorem 4.1.

Theorem 4.2 *For any dynamic scenario in the leaf-increasing tree model, where the initial number of nodes in the tree is n_0 and n is the final number of nodes in the tree, $\text{SDL}^{k(x)}$ is a dynamic F -labeling scheme, satisfying the following complexities.*

- $\mathcal{LS}(\text{SDL}^{k(x)}, n) = O(\log_{k(n)} n \cdot \mathcal{LS}(\pi, n))$.
- $\mathcal{MC}(\text{SDL}^{k(x)}, n) = O(k(n)(\log_{k(n)} n)\mathcal{MC}(\pi, n))$.

By examining the details in [3, 13, 31, 30, 22] concerning the labeling schemes supporting the above mentioned functions (i.e., the ancestry relation, the label-based and the id-based NCA relations, the separation level, the distance and the routing functions), it can be easily shown that for each of the above mentioned labeling schemes π , there exists a distributed protocol assigning the labels of π on static trees using a linear number of messages. Therefore, by setting $k(n) = \log^\epsilon n$ for any $\epsilon > 0$, we obtain the following corollary.

Corollary 4.3 *In the leaf-increasing tree model, there exist dynamic labeling schemes with message complexity $O(n \frac{\log^{1+\epsilon} n}{\log \log n})$ for the following functions.*

- *For the routing function in the designer port model, the ancestry relation and the label-based NCA relation: with label size $O(\frac{\log^2 n}{\log \log n})$.*
- *For the routing function in the adversary port model: with label size $O(\frac{\log^3 n}{\log^2 \log n})$.*
- *For the distance function, the separation level and the id-based NCA relation: with label size $O(\frac{\log^3 n}{\log \log n})$.*

4.2 The dynamic F -labeling scheme $DL^{k(x)}$

In the leaf-dynamic tree model, each vertex u may store information in $Memory(u)$ that is required for correct performances of our dynamic schemes. One of the difficulties that may rise is that when a leaf u is deleted, we lose the information stored in u . In order to overcome this difficulty, we use the following backup procedure. Throughout the dynamic scenario we maintain for every child u of a non-leaf node v , a copy of $Memory(u)$ stored as backup in either v or in a sibling of u . Thus, when u is deleted, v retrieves the information in $Memory(u)$ by communicating with the vertex holding the corresponding copy. It turns out that the above mentioned backup procedure can be implemented without increasing the asymptotic label size and message complexity. The details regarding this implementation appear in [21].

Before turning to describe our main dynamic labeling scheme for the leaf-dynamic model, let us first describe a version of Scheme $SDL^{k(x)}$, denoted $SEM-DL^{k(x)}$, which operates in the leaf-dynamic tree model and mimics the behavior of Scheme $SDL^{k(x)}$ on the dynamic scenario assuming deletions are never made. Recall that in the leaf-increasing tree model, Scheme $SDL^{k(x)}$ occasionally invokes sub-protocol $RESET$ on different subtrees T' and that in the first step of this sub-protocol, the current number of nodes in T' is calculated. In the leaf-dynamic tree model, Scheme $SEM-DL^{k(x)}$ carries out the same steps as $SDL^{k(x)}$ except for the following two modifications.

- 1) Messages are not passed to deleted vertices.
- 2) Every time Scheme $SDL^{k(x)}$ invokes Sub-protocol $RESET(T')$, instead of calculating the current number of nodes in T' in the first step of Sub-protocol $RESET(T')$, Scheme $SEM-DL^{k(x)}$ calculates the number of nodes that have ever been in T' , i.e., the existing nodes in T' together with the

deleted ones.

The first modification is implemented trivially. Let us now describe how to implement the second modification. Recall that by the subtrees decomposition properties, each vertex v belongs to precisely one l -level subtree for each $1 \leq l \leq p$. Moreover, an l -global reset operation is invoked only on l -level subtrees. At any given time, let T_l be some l -level subtree. Let $n^-(T_l)$ denote the number of nodes that have been deleted from T_l and let $n(T_l) = |T_l| + n^-(T_l)$, i.e., the number of nodes that have ever been in T_l . Throughout the dynamic scenario, for every $1 \leq l \leq p$, each vertex v keeps a counter $\omega_l(v)$ such that the following invariant is maintained at all times for every l -level subtree T_l .

The T_l -invariant: $\sum_{\{v \in T_l\}} \omega_l(v) = n(T_l)$.

Assuming that for every l -level subtree, the T_l -invariant holds at all times, we now show how to implement the second modification. Instead of calculating the current number of nodes in T_l in the first step of Sub-protocol $\text{RESET}(T_l)$, we calculate $n(T_l)$ using broadcast and upcast operations (see [32]) on T_l by which $\sum_{\{v \in T_l\}} \omega_l(v)$ is calculated.

We now describe how Scheme $\text{SEM-DL}^{k(x)}$ guarantees that the T_l -invariant is maintained for every subtree T_l .

1. If Scheme $\text{SEM-DL}^{k(x)}$ is invoked on the initial tree T_0 then let p such that FSDL_p^k is initially invoked on T_0 . For every vertex $v \in T_0$, set $\omega_l(v) = 1$ for every $1 \leq l \leq p$.
2. If v is added as a leaf to the tree then v sets $\omega_l(v) = 1$ for every $1 \leq l \leq p$.
3. If v participates in some Sub-protocol RESET which is invoked by some Scheme $\text{FSDL}_{l'}^k$ then for every $1 \leq l < l'$, v sets $\omega_l(v) = 1$.
4. If a child u of v is deleted, then v extracts $\{\omega_l(u) \mid 1 \leq l \leq p\}$ using the copy of $\text{Memory}(u)$ (as mentioned before). Subsequently, v sets $\omega_l(v) = \omega_l(v) + \omega_l(u)$ for every $1 \leq l \leq p$ such that $T_l(v) = T_l(u)$.

Using induction on the time, it is easy to verify that for every subtree T_l , the T_l -invariant is indeed maintained at all times. Therefore, Scheme $\text{SEM-DL}^{k(x)}$ can implement the modifications to Scheme $\text{SDL}^{k(x)}$ described above. Thus, using the same steps as in the proof of Theorem 4.2, we obtain the following lemma.

Lemma 4.4 *For any dynamic scenario in the leaf-dynamic tree model, where the initial number of nodes in the tree is n_0 and n^+ additions are made, Scheme $\text{SEM-DL}^{k(x)}$ is a dynamic F -labeling scheme with the following complexities. Let $n' = n_0 + n^+$.*

- $\mathcal{LS}(\text{SEM-DL}^{k(x)}, n) = O(\log_{k(n')} n' \cdot \mathcal{LS}(\pi, n'))$.

- $\mathcal{MC}(\text{SEM-DL}^{k(x)}, n') = O(k(n')(\log_{k(n')} n')\mathcal{MC}(\pi, n'))$.

We now turn to describe Scheme $\text{DL}^{k(x)}$ which is designed to operate in the leaf-dynamic tree model and improves the complexities of $\text{SEM-DL}^{k(x)}$. Scheme $\text{DL}^{k(x)}$ uses a method similar to the one presented in Subsection 3.4 in [29]. The general idea is to run, in parallel to $\text{SEM-DL}^{k(x)}$, a protocol for estimating the number of topological changes in the tree. Every $\Theta(n)$ topological changes we restart Protocol $\text{SEM-DL}^{k(x)}$ again on the current initial tree T_0 .

Denote by τ the number of topological changes made to the tree during the execution in the leaf-dynamic tree model. Fix $\delta = 9/8$. We use Protocol CHANGEWATCH from [29] (which is an instance of the protocol of [1]) in which the root maintains an estimate $\tilde{\tau}$ of τ . This is done by applying the same mechanism as in Protocol WEIGHTWATCH from [29] separately for the additions of vertices and for the deletions. I.e, we run two protocols in parallel. The first is designed to count the additions. In order to do that, we ignore the deletions and perform the same steps as in Protocol WEIGHTWATCH . The second protocol is designed to count the deletions. For this we ignore the additions, and carry the same steps as in Protocol WEIGHTWATCH , except for deletions rather than for additions. Let n_0 be the number of vertices in the tree when Protocol CHANGEWATCH was initiated. Let n_+ and n_- be the number of additions and deletions respectively and let \tilde{n}_+ and \tilde{n}_- be the root's estimated number of additions and deletions respectively.

As mentioned in Section 3.4.1 of [29], as long as the root's estimates satisfy $\tilde{n}_+ \leq \frac{n_0}{9}$ and $\tilde{n}_- \leq \frac{n_0}{9}$, it is guaranteed that $\tau = n_+ + n_- \leq \frac{n_0}{2}$. Moreover, if $\tilde{n}_+ > \frac{n_0}{9}$ or $\tilde{n}_- > \frac{n_0}{9}$ then $\tau > \frac{n_0}{9}$. As mentioned in [29], $\mathcal{MC}(\text{CHANGEWATCH}, \bar{n}) = O(\sum_i \log^2 n_i)$.

Protocol $\text{DL}^{k(x)}$ operates as follows.

Scheme $\text{DL}^{k(x)}$

1. Let T_0 be the current tree. The root initiates a convergecast process in order to calculate n_0 , the initial number of nodes in the tree.
2. Protocols $\text{SEM-DL}^{k(x)}$ and CHANGEWATCH are started on T_0 .
3. When one of the estimates \tilde{n}_+ or \tilde{n}_- exceeds $n_0/9$, return to Step 1.

Theorem 4.5 $\text{DL}^{k(x)}$ is a dynamic F -labeling scheme for the leaf-dynamic tree model, satisfying the following properties.

- $\mathcal{LS}(\text{DL}^{k(x)}, n) = O(\log_{k(n)} n \cdot \mathcal{LS}(\pi, n))$.
- $\mathcal{MC}(\text{DL}^{k(x)}, \bar{n}) = O\left(\sum_i k(n_i)(\log_{k(n_i)} n) \frac{\mathcal{MC}(\pi, n_i)}{n_i}\right) + O(\sum_i \log^2 n_i)$.

Proof: Scheme $\text{DL}^{k(x)}$ is restarted by returning to Step 1 after τ topological changes, for $\tau = \Theta(n_0)$, where n_0 is the last recorded tree size at Step 1. Consequently, the current tree size satisfies

$n = \Theta(n_0)$ and $n = \Theta(n_0 + n^+)$ where n^+ is the number of additions made from the last time Step 1 was invoked. Therefore, by the first part of Lemma 4.4 and by our assumptions on $k(n)$ and $\mathcal{LS}(\pi, n)$, we obtain the first part of the theorem.

Let us now turn to prove the second part of the theorem. Let i_1, \dots, i_m be the indices of the topological changes on which Scheme $\text{DL}^{k(x)}$ returns to Step 1. Denote by M_l the number of messages resulting from the l 'th time until the $l + 1$ 'st time Scheme $\text{SEM-DL}^{k(x)}$ is applied in Step 2 of Scheme $\text{DL}^{k(x)}$. Clearly

$$\mathcal{MC}(\text{DL}^{k(x)}, \bar{n}) = \sum_{l=1}^m M_l + \mathcal{MC}(\text{CHANGEWATCH}, \bar{n}) .$$

Since the number of changes relevant to M_l is $\Theta(n_{i_l})$, by our assumptions on $k(\cdot)$ and $\mathcal{MC}(\pi, \cdot)$, we obtain

$$M_l \leq O\left(k(n_{i_l})(\log_{k(n_{i_l})} n) \mathcal{MC}(\pi, n_{i_l})\right) .$$

Again, by our assumptions on $\mathcal{MC}(\pi, \cdot)$ and $k(\cdot)$ and we actually have that

$$M_l \leq O\left(\sum_{j=i_l}^{i_{l+1}-1} k(n_j)(\log_{k(n_j)} n) \frac{\mathcal{MC}(\pi, n_j)}{n_j}\right) .$$

Therefore

$$\sum_{l=1}^m M_l \leq O\left(\sum_j k(n_j)(\log_{k(n_j)} n) \frac{\mathcal{MC}(\pi, n_j)}{n_j}\right) .$$

Since $\mathcal{MC}(\text{CHANGEWATCH}, \bar{n}) = O(\sum_i \log^2 n_i)$, the second part of the theorem follows. \blacksquare

By setting $k(x) = n^\epsilon$ for any $0 < \epsilon < 1$, we obtain the following corollary.

Corollary 4.6 • *In the leaf-dynamic tree model, for every static F -labeling scheme π , there exists a dynamic F -labeling scheme with the same asymptotic label size as π and sublinear amortized message complexity.*

- *In the leaf-dynamic tree model, there exist dynamic labeling schemes for the ancestry relation, the id-based and label-based NCA relations and for the routing function (both in the designer and the adversary port models) using asymptotically optimal label sizes and sublinear amortized message complexity.*

By setting $k(x) = \log^\epsilon n$ for any $0 < \epsilon < 1$, we obtain the following corollary.

Corollary 4.7 • *In the leaf-dynamic tree model, for every static F -labeling scheme π , there exists a dynamic F -labeling scheme with $O\left(\sum_i \frac{\log^{1+\epsilon} n_i}{\log \log n_i} \cdot \frac{\mathcal{MC}(\pi, n_i)}{n_i}\right) + O(\sum_i \log^2 n_i)$ message complexity and multiplicative overhead of $O\left(\frac{\log n}{\log \log n}\right)$ over the label size of π .*

- *In the leaf-dynamic tree model, there exist dynamic labeling schemes for the ancestry relation, the id-based and label-based NCA relations and for the routing function (both in the designer*

and the adversary port models), with message complexity $O(\sum_i \log^2 n_i)$ and multiplicative overhead of $O(\frac{\log n}{\log \log n})$ over the corresponding asymptotically optimal label size.

5 External memory complexity

5.1 Types of memory

We distinguish between three types of memory bits used by a node v . The first type consists of the bits in the label $\mathcal{M}(v)$ given to v by the marker algorithm. The second type consists of the memory bits used by the static algorithm π in order to calculate the static labels. The third type of bits, referred to as the *external memory* bits, consists of the additional external storage used during updates and maintenance operations by the dynamic labeling scheme. As mentioned before, for certain applications (and particularly routing), the label $\mathcal{M}(v)$ seems to be a more critical consideration than the total amount of storage needed for the information maintenance. In addition, the second type of memory bits are used by the static algorithm π only when it is invoked, which is done infrequently. Moreover, we note that by examining the details in [3, 13, 31, 30, 22] concerning the labeling schemes supporting the ancestry relation, the label-based and the id-based NCA relations, as well as the separation level, distance and routing functions, it can be easily shown that for each of the above mentioned labeling schemes π , there exists a distributed protocol assigning the labels of π on static trees using a linear number of messages. Moreover, at any vertex, the number of memory bits used by these static algorithms is asymptotically the same as the label size.

In the following discussion, we therefore try to minimize the number of external memory bits used by our dynamic schemes. Let us first describe the need for these memory bits.

Consider either Scheme $\text{DL}^{k(x)}$ or Scheme $\text{SDL}^{k(x)}$ for some function $k(x)$. Recall that at any time during the dynamic scenario, there exists parameters k and p such that the only FSDL schemes that are currently invoked are of the form FSDL_l^k where $1 \leq l \leq p$. Moreover, every vertex v belongs to precisely one l -level subtree, namely $T_l(v)$, for each $1 \leq l \leq p$. Therefore, each node v holds at most p counters of the form μ_l and at most p counters of the form $\omega_l(v)$. Since each such counter contains $O(\log n)$ bits we get that holding these counters incurs $O(\log_{k(n)} n \cdot \log n)$ external memory bits per node.

Since each node v may participate at the same time in different schemes FSDL_l^k for different l 's, v must know, for each $1 \leq l \leq p$, which of its edges correspond to $T_l(v)$, its l -level subtree. Naively storing this information at v may incur $\Omega(p \cdot n)$ bits of memory. Note that in Scheme FSDL_l^k , each vertex v either communicates with its parent in $T_l(v)$ or with all its children in $T_l(v)$. Moreover, for each l , if v is not the root of $T_l(v)$ then its parent in $T_l(v)$ is its parent in T , namely, $\text{parent}(v)$. Therefore, in order for v to know, for each l , which of its ports leads to its parent in $T_l(v)$, it is

enough for it to know which port leads to $\text{parent}(v)$ and for each l to keep a bit, indicating whether v is the root of $T_l(v)$ or not. This costs $O(p + \log n)$ memory bits. For each $1 \leq l \leq p$, let $E_l(v)$ be the port numbers (at v) corresponding to the edges connecting v to its children in $T_l(v)$. Note that E_p is precisely the collection of all port numbers at v leading to v 's children in T . It is therefore enough to ensure that v is able to detect, for every $1 \leq l < p$, which of its port numbers are in $E_l(v)$.

We first consider our schemes in the designer port model, and then discuss them in the adversary port model. In the designer port model we show that the external memory bits used by a vertex do not exceed the asymptotic label size of the corresponding dynamic scheme. However, in the adversary port model, for a given static scheme π , if the port numbers given by the adversary use many bits (in comparison to the the label size of π), then the external memory bits used by a vertex may exceed the asymptotic label size of the corresponding dynamic scheme. Let us note that assuming the designer port model, if the port numbers are encoded in the labels of the corresponding static labeling scheme (e.g. the routing scheme of [13] for the designer port model) then we cannot re-enumerate the port numbers to save external memory bits. In the context of this section, we therefore consider such schemes as operating in the adversary port model. Let us note, however, that in the designer port model, of all the above mentioned functions F , the only known static F -labeling scheme which actually encodes the port numbers into the labels, is the routing scheme of [13] (for the designer port model). Since this static routing labeling scheme uses port numbers which are encoded using only $O(\log n)$ bits, the external memory complexity of our corresponding dynamic routing labeling schemes is asymptotically the same as the label size. (See Corollary 5.3).

5.2 External memory in the designer port model

In the designer port model, in order to reduce the memory storage used at each node, we exploit the fact that the tree $T_l(v)$ is a subtree of $T_{l+1}(v)$. This is done in the following manner. For every $1 \leq l < p$, each node v keeps a variable a_l and maintains an enumeration of its ports so that the following invariant is maintained at all times.

The designer-invariant: For every $l = 1, 2, \dots, p - 1$, $\{1, 2 \dots, a_l\} = E_l(v)$. In other words, v maintains an enumeration of its ports so that the port numbers from 1 to a_l correspond to the edges connecting v to its children in $T_l(v)$.

Using the designer-invariant, the port numbers in $E_l(v)$ can easily be identified by v since they are precisely the port numbers $1, 2, \dots, a_l$. The details regarding the implementation of the designer invariant appear in [21]. We note that this implementation does not affect the asymptotic complexities of our schemes.

Since each vertex v holds $O(p)$ counters and at most p variables of the form a_l , and since each

of these variables and counters contains $O(\log n)$ bits, we obtain the following lemma.

Lemma 5.1 *In the designer port model, for any execution of either Scheme $\text{SDL}^{k(x)}$ in the leaf-increasing tree model or Scheme $\text{DL}^{k(x)}$ in the leaf-dynamic tree model, the maximal number of external memory bits used by a vertex in any n -node tree is $O(\log_{k(n)} n \cdot \log n)$.*

5.3 External memory in the adversary port model

We first remark that in [29], the designer port model is assumed. Since port numbers are used in the labels given by the dynamic schemes of [29], applying their scheme in the adversary port model may affect the label sizes of the schemes. Specifically, let $\tau(n)$ be the maximum port number given by the adversary to any node in any n -node tree, taken over all scenarios. Then the upper bound on the label sizes of the general schemes proposed in [29] changes from $O(d \log_d n \cdot \mathcal{LS}(\pi, n))$ to $O(d \log_d n \cdot (\mathcal{LS}(\pi, n) + \log \tau(n)))$ (see Lemma 4.12 in [29]). In contrast, applying our schemes in the adversary port model may only affect the external memory complexities. As discussed before, it is enough to guarantee that each node v knows for each $l < p$ which of its port numbers is in $E_l(v)$. Informally, in order to implement the above, each node v distributes the relevant information to its children in T and collects it back when needed. The proofs of the following lemmas are quiet technical and they appear in [21].

Lemma 5.2 *Assuming the adversary port model, in either Scheme $\text{SDL}^{k(x)}$ or Scheme $\text{DL}^{k(x)}$, the maximal number of external memory bits used by a vertex is $O(\log_{k(n)} n \cdot (\log \tau(n) + \log n))$.*

As mentioned before, in the designer port model, if the corresponding static labeling scheme encodes the port numbers into the labels, in the context of saving external memory bits, we consider such a scheme as operating in the adversary port model. However, of all the above mentioned functions F , the only known such static F -labeling scheme is the routing labeling scheme of [13] (for the designer port model). Since this static scheme uses port numbers which are encoded using $O(\log n)$ bits, we obtain the following corollary.

Corollary 5.3 *Let π be the static routing scheme of [13] for the designer port model. Let $\hat{\pi}$ be either Scheme $\text{SDL}^{k(x)}$ or Scheme $\text{DL}^{k(x)}$ extending π to the dynamic scenario. Then the maximal number of external memory bits used by a vertex in $\hat{\pi}$ is $O(\log_{k(n)} n \cdot \log n)$ (which is asymptotically the same as the label size of $\hat{\pi}$).*

References

- [1] Y. Afek, B. Awerbuch, S.A. Plotkin and M. Saks. Local management of a global resource in a communication. *J. of the ACM* **43**, (1996), 1–19.
- [2] S. Alstrup, P. Bille and T. Rauhe. Labeling schemes for small distances in trees. In *Proc. 14th ACM-SIAM Symp. on Discrete Algorithms*, Jan. 2003.

- [3] S. Alstrup, C. Gavoille, H. Kaplan and T. Rauhe. Nearest Common Ancestors: A Survey and a new Distributed Algorithm. *Theory of Computing Systems* **37**, (2004), 441–456.
- [4] S. Alstrup, J. Holm and M. Thorup. Maintaining Center and Median in Dynamic Trees. In *Proc. 7th Scandinavian Workshop on Algorithm Theory*, July. 2000.
- [5] S. Abiteboul, H. Kaplan and T. Milo. Compact labeling schemes for ancestor queries. In *Proc. 12th ACM-SIAM Symp. on Discrete Algorithms*, Jan. 2001.
- [6] S. Alstrup and T. Rauhe. Improved Labeling Scheme for Ancestor Queries. In *Proc. 19th ACM-SIAM Symp. on Discrete Algorithms*, Jan. 2002.
- [7] M.A. Breuer and J. Folkman. An unexpected result on coding the vertices of a graph. *J. of Mathematical Analysis and Applications* **20**, (1967), 583–600.
- [8] M.A. Breuer. Coding the vertexes of a graph. *IEEE Trans. on Information Theory*, IT-12:148–153, 1966.
- [9] R. Cole and R. Hariharan. Dynamic LCA Queries on Trees. In *Proc. 10th ACM-SIAM Symp. on Discrete Algorithms*, 235–244, 1999.
- [10] E. Cohen, E. Halperin, H. Kaplan and U. Zwick. Reachability and Distance Queries via 2-hop Labels. In *Proc. 13th ACM-SIAM Symp. on Discrete Algorithms*, Jan. 2002.
- [11] E. Cohen, H. Kaplan and T. Milo. Labeling dynamic XML trees. In *Proc. 21st ACM Symp. on Principles of Database Systems*, June 2002.
- [12] D. Eppstein, Z. Galil and G. F. Italiano. Dynamic Graph Algorithms. In *Algorithms and Theoretical Computing Handbook*, M. J. Atallah, ed., CRC Press, 1999, chapter 8.
- [13] P. Fraigniaud and C. Gavoille. Routing in trees. In *Proc. 28th Int. Colloq. on Automata, Languages & Prog.*, LNCS 2076, pages 757–772, July 2001.
- [14] P. Fraigniaud and C. Gavoille. A space lower bound for routing in trees. In *Proc. 19th Symp. on Theoretical Aspects of Computer Science*, Mar. 2002.
- [15] J. Feigenbaum and S. Kannan. Dynamic Graph Algorithms. In *Handbook of Discrete and Combinatorial Mathematics*, CRC Press, 2000.
- [16] C. Gavoille and C. Paul. Split decomposition and distance labelling: an optimal scheme for distance hereditary graphs. In *Proc. European Conf. on Combinatorics, Graph Theory and Applications*, Sept. 2001.
- [17] C. Gavoille and D. Peleg. Compact and Localized Distributed Data Structures. *J. of Distributed Computing* **16**, (2003), 111–120.
- [18] C. Gavoille, M. Katz, N.A. Katz, C. Paul and D. Peleg. Approximate Distance Labeling Schemes. In *9th European Symp. on Algorithms*, Aug. 2001, Aarhus, Denmark, SV-LNCS 2161, 476–488.
- [19] C. Gavoille, D. Peleg, S. Pérennes and R. Raz. Distance labeling in graphs. In *Proc. 12th ACM-SIAM Symp. on Discrete Algorithms*, pages 210–219, Jan. 2001.
- [20] J. Holm, K. Lichtenberg and M. Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *J. of the ACM* **48(4)**, (2001), 723–760.

- [21] A. Korman. General Compact Labeling Schemes for Dynamic Trees. *ACM Comp. Research Repository* <http://arxiv.org/abs/cs.DC/0605141>.
- [22] S. Kannan, M. Naor, and S. Rudich. Implicit representation of graphs. In *SIAM J. on Discrete Math* **5**, (1992), 596–603.
- [23] H. Kaplan and T. Milo. Short and simple labels for small distances and other functions. In *Workshop on Algorithms and Data Structures*, Aug. 2001.
- [24] H. Kaplan and T. Milo. Parent and ancestor queries using a compact index. In *Proc. 20th ACM Symp. on Principles of Database Systems*, May 2001.
- [25] H. Kaplan, T. Milo and R. Shabo. A Comparison of Labeling Schemes for Ancestor Queries. In *Proc. 19th ACM-SIAM Symp. on Discrete Algorithms*, Jan. 2002.
- [26] M. Katz, N.A. Katz, A. Korman and D. Peleg. Labeling schemes for flow and connectivity. In *Proc. 19th ACM-SIAM Symp. on Discrete Algorithms*, Jan. 2002.
- [27] M. Katz, N.A. Katz, and D. Peleg. Distance labeling schemes for well-separated graph classes. In *Proc. 17th Symp. on Theoretical Aspects of Computer Science*, pages 516–528, February 2000.
- [28] A. Korman and D. Peleg. Labeling Schemes for Weighted Dynamic Trees. In *Proc. 30th Int. Colloq. on Automata, Languages & Prog.*, Eindhoven, The Netherlands, July 2003, Springer LNCS.
- [29] A. Korman, D. Peleg and Y. Rodeh. Labeling schemes for dynamic tree networks. *Theory of Computing Systems* **37**, (2004), 49–75.
- [30] D. Peleg. Proximity-preserving labeling schemes and their applications. In *Proc. 25th Int. Workshop on Graph-Theoretic Concepts in Computer Science*, pages 30–41, June 1999.
- [31] D. Peleg. Informative labeling schemes for graphs. In *Proc. 25th Symp. on Mathematical Foundations of Computer Science*, volume LNCS-1893, pages 579–588. Springer-Verlag, Aug. 2000.
- [32] D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, 2000.
- [33] D. D. Sleator and R. E. Tarjan. A data structure for dynamic trees. *Journal of Computer and System Sciences* **26(1)**, (1983), 362–391.
- [34] N. Santoro and R. Khatib. Labelling and implicit routing in networks. *The Computer Journal* **28**, (1985), 5–8.
- [35] M. Thorup. Compact oracles for reachability and approximate distances in planar digraphs. *J. of the ACM* **51**, (2004), 993–1024.
- [36] M. Thorup and U. Zwick. Compact routing schemes. In *Proc. 13th ACM Symp. on Parallel Algorithms and Architecture*, pages 1–10, Hersonissos, Crete, Greece, July 2001.