Combine and Conquer

Robert F. Cohen and Roberto Tamassia

Department of Computer Science
Brown University
Providence, Rhode Island 02912

CS-92-19
April 1992

Combine and Conquer *

Robert F. Cohen
Roberto Tamassia
Department of Computer Science
Brown University
Providence, Rhode Island 02912–1910
(rfc@cs.brown.edu, rt@cs.brown.edu)

Abstract

We present a general technique for dynamizing a class of problems whose underlying structure is a computation graph embedded in a tree. We associate values, called *attributes*, with the nodes, paths, and subtrees of our trees. Path attributes form a *path attribute system*, if they are maintained in constant time under path concatenation. Additionally, attributes form a *tree attribute system* if the tree attributes of the tail of a path Π are determined in constant time from the path attributes of Π.

We also introduce a new data structure called a *linear attribute grammar*. An *attribute grammar* is a tree-based expression where the values a node μ are calculated from the values at the parent, siblings, and/or the children of μ. A linear attribute grammar, is an attribute grammar where all dependencies are linear.

Our contributions can be summarized as follows. We provide a framework for maintaining attribute systems on trees in a fully dynamic environment. We show that given a semiring $\mathcal{S}$, a set of linear expressions with binary and unary operators over $\mathcal{S}^k$ can be dynamically maintained in a fully dynamic environment using linear space and logarithmic time per operation. We show that a linear attribute grammar can be dynamically maintained in a fully dynamic environment using linear space and logarithmic time per operation. Finally, we present many examples of application of our technique to dynamic graph and geometric problems.

(April 1992)

*Research supported in part by the National Science Foundation under grant CCR-9007851, by the U.S. Army Research Office under grant DAAL03-91-G-0035, and by the Office of Naval Research and the Defense Advanced Research Projects Agency under contract N00014-91-J-4052, ARPA order 8225.

1 Introduction

The development of dynamic algorithms and data structures is a challenging area of research that has attracted increasing interest in the last years. While considerable progress has been made in geometric searching problems (see, e.g., the survey paper [5]), considerably fewer results exist on the dynamization of graph algorithms. A crucial development in the area of dynamic geometric searching has been the identification of general methods that apply to a large class of problems. In particular, the techniques developed by Overmars et al. (summarized in [21]) for the class of decomposable search problems constitute a fundamental contribution toward the dynamization of large classes of geometric problems. The availability of such general techniques appears instead to be lacking in the area of dynamic graph algorithms. The goal of this research is exactly to provide such generalized techniques in the realm of dynamic tree and graph problems.

Our approach is motivated by the observation that a number of dynamic graph algorithms [10,11,13,14,15,16,19,23,31], developed mostly for connectivity problems, appear to be based on the following fundamental idea: Decompose a graph into subgraphs with limited overlap, and represent such a decomposition by means of a tree so that dynamic operations on the graph are reflected into corresponding dynamic tree operations, which are in turn supported by variations of the link-cut trees of Sleator and Tarjan [28].

We associate values, called *attributes*, with the nodes, paths, and subtrees of our trees. Path attributes form a *path attribute system*, if they are maintained in constant time under path concatenation. Additionally, attributes form a *tree attribute system* if the tree attributes of the tail of a path Π are determined in constant time from the path attributes of Π.

We also introduce a new data structure called a *linear attribute grammar*. An *attribute grammar* is a tree-based expression where the values a node μ are calculated from the values at the parent, siblings, and/or the children of μ. A linear attribute grammar, is an attribute grammar where all dependencies are linear.

Our contributions can be summarized as follows:

- We provide a framework for maintaining attribute systems on trees in a fully dynamic environment. Our technique extends and generalizes the dynamic trees of [28] and the dynamic expression trees of [9]. Our technique also extends and generalizes also the work on order-decomposable search problems by Overmars [Overmars Notes 83].
- We show that given a semiring $\mathcal{S}$, a set of linear expressions with binary and unary operators over $\mathcal{S}^k$ can be dynamically maintained in a fully dynamic environment using linear space and logarithmic time per operation.
- We show that a linear attribute grammar can be dynamically maintained in a fully dynamic environment using linear space and logarithmic time per operation. Linear attribute grammars can be used as the data structure for dynamic algorithms for several problems in graph drawing.

The rest of this paper is organized as follows. Section 2 describes fully dynamic algorithms to maintain attributes on paths and trees. Separate algorithms are presented for trees of bounded and unbounded degrees. Section 3.1 applies the results of the previous section to present a fully dynamic algorithm for maintaining the solutions of linear expressions. Section 3.2 presents a fully dynamic algorithm for linear attribute grammars. Finally, section 4

presents other applications including two types of generalized heaps.

2 Attribute Systems on Paths and Trees

In this section we present a framework for dynamically maintaining an important class of path and tree attributes. We introduce the concept of path attribute systems and tree attribute systems which significantly extend and generalize the dynamic trees of [28]. We begin by discussing dynamic algorithms on a collection of paths, and then show that trees can be maintained as a collection of paths.

2.1 Path Attribute Systems

Paths are directed. The first and last nodes of a path Π are called the *head* and *tail* of Π, respectively, and are denoted with $head(\Pi)$ and $tail(\Pi)$. The *reversed path* $\overline{\Pi}$ of Π is the path obtained by reversing all edges of Π. Formally, we have:

Definition 1 *A* path *is a collection of nodes and directed edges such that:*

- *The empty path with no nodes is a path.*
- *A single node μ is a path with* $\text{head}(\mu) = \text{tail}(\mu) = \mu$.
- *If Π' and Π'' are paths then the concatenation of Π' and Π'' is a path Π formed by adding a directed edge from* $\text{tail}(\Pi')$ *to* $\text{head}(\Pi'')$.

Definition 2 *A* node attribute *N is a function on nodes. The values N can take are arbitrary, but can be stored in $O(1)$ space. A* path attribute *P is a function on paths. The value of $P(\Pi)$ for a path Π is dependent on the values of $N(\mu)$ for each node μ on Π, and on the order of the nodes of Π. The value of $P(\Pi)$ can be stored in $O(1)$ space.*

Definition 3 *If the values of a node attribute N are taken from a monoid, then N is said to be* globally updatable, *otherwise N is* locally updatable. *We define $\mathcal{N}_G(\mu)$ and $\mathcal{N}_L(\mu)$ to be the set of globally and locally updatable node attributes of μ.*

Definition 4 *A* node attribute set *$\mathcal{N}$ is a finite collection of node attributes $N_1, \ldots, N_r$. Similarly, a* path attribute set *$\mathcal{P}$ is a finite collection of path attributes $P_1, \ldots, P_s$. For a node μ, the* value *of $\mathcal{N}(\mu)$ is the vector $(N_1(\mu), \ldots, N_r(\mu))$. For a path Π, the value of $\mathcal{P}(\Pi)$ is the vector $(P_1(\Pi), \ldots, P_s(\Pi))$.*

For each path Π, we consider the node attribute sets $\mathcal{N}(head(\Pi))$ and $\mathcal{N}(tail(\Pi))$ to be included in path attribute set $\mathcal{P}(\Pi)$.

Definition 5 *Suppose $\mathcal{N}$ is a node attribute set $\mathcal{P}$ is a path attribute set, and $F : \mathcal{P} \times \mathcal{P} \to \mathcal{P}$ is a function . The triple $(\mathcal{N}, \mathcal{P}, F)$ is a* path attribute system *$\mathcal{Q}$ if for any path Π that is the concatenation of paths Π' and Π'', $\mathcal{P}(\Pi)$ can be determined in $O(1)$ time as $F(\mathcal{P}(\Pi'), \mathcal{P}(\Pi''))$. Function F is called the* concatenation function *of $\mathcal{Q}$.*

Example 1 *Suppose for each node μ we keep the following node attribute:*

- $\text{weight}(\mu)$ — *the real-valued weight of node μ.*

Additionally, we have the following path attributes for a path Π:

- $\mathrm{sum}(\Pi)$ — *the sum of the weights of the nodes of* Π.
- $\mathrm{squaredsum}(\Pi)$ — *the square of the sum of the weights of the nodes of* Π.

Suppose path Π *is the concatenation of paths* Π' *and* Π''. *If our path attribute set* $\mathcal{P}(\Pi)$ *consists solely of* $\mathrm{squaredsum}(\Pi)$, *then we can not find a concatenation function to maintain* $\mathcal{P}(\Pi)$, *since we cannot determine* $\mathrm{squaredsum}(\Pi)$ *from* $\mathrm{squaredsum}(\Pi')$ *and* $\mathrm{squaredsum}(\Pi'')$. *However, if* $\mathcal{P}(\Pi)$ *is the pair* $(\mathrm{sum}(\Pi), \mathrm{squaredsum}(\Pi))$, *then we can use concatenation function* F *such that:*

$$\begin{aligned} \mathrm{sum}(\Pi) &= \mathrm{sum}(\Pi') + \mathrm{sum}(\Pi'') \\ \mathrm{squaredsum}(\Pi) &= (\mathrm{sum}(\Pi))^2 \end{aligned}$$

Each node attribute N has a *secondary value* $N^*(\mu)$ in addition to its "normal" *primary value* $N(\mu)$. For node attribute set $\mathcal{N} = (N_1, \ldots, N_r)$, $\mathcal{N}^*$ denotes $(N_1^*, \ldots, N_r^*)$. We also have corresponding secondary values for path attributes, and similarly define $\mathcal{P}^*(\Pi)$ for a path Π.

A decomposable search problem π [21] on a set D, locates a distinguished element $x = \pi(D)$ with the conditions of given any partition of D into subsets D' and D'', we can determine in constant time if x is in D' or D''. A number of methods are presented to maintain the solution for a decomposable search problem.

We extend this notion to paths as follows:

Definition 6 *Given a path attribute system* $\mathcal{Q} = (\mathcal{N}, \mathcal{P}, F)$, *a* path-selection query Q *for* $\mathcal{Q}$ *maps a path* Π *and a query argument* q *into a node* $\mu = Q(\Pi, q)$ *of* Π. *Suppose path* Π *is the concatenation of* Π' *and* Π''. *A* path-selection function $S(\Pi, q)$ *for* Q, *determines in* $O(1)$ *time whether* μ *is in* Π' *or* Π'' *from* q *and the values* $\mathcal{P}(\Pi')$ *and* $\mathcal{P}(\Pi'')$.

We will show that we can perform a selection query given the associated selection function.

We consider a rather general dynamic environment for a set of paths equipped with a path attribute system $\mathcal{Q}$ and a collection $\mathcal{S}$ of path-selection functions. Update operations include: *splitting*, *concatenating*, and *reversing paths*; and *updating node attributes*; and *exchanging* primary and secondary attribute values for all the nodes on a path. Query operations include: *evaluating* node and path attributes, and *computing path-selection functions*.

We allow both "local" and "global" updates of node attributes. A *local update* consists of changing the value $N(\mu)$ of a single node μ. A *global update* can be performed on a globally updatable node attribute N with values taken from a monoid with operator $\odot$, and consists of applying an incremental change δ to the attribute of every node in a path Π. Formally, a global update sets $N(\mu) = N(\mu) \odot \delta$ for every node μ of Π, where we assume that $x \odot y$ can be computed in $O(1)$ time. If a path attribute P depends on a globally updatable node attribute N, we require that the new value of $P(\Pi)$ can be computed in $O(1)$ time given the variation δ of N in Π.

We support the following repertory of operations on a set of (uniquely identified) paths equipped with a path attribute system $\mathcal{Q} = (\mathcal{N}, \mathcal{P}, F)$ and a collection $\mathcal{F}$ of path-selection functions for $\mathcal{Q}$:

- *makepath*(*node* ν; *value* x) — return a new path consisting of a single node μ and set $\mathcal{N}(\nu) = x$.
- *deletepath*(*node* ν) — remove the single-node path consisting of node ν.

- *getpath*(*node* ν) — return the unique identifier of the path containing ν.
- *head*(*path* Π) — return the head of path Π.
- *tail*(*path* Π) — return the tail of path Π.
- *before*(*node* μ, *integer* k) — return the k nodes preceding μ on *getpath*(μ). If there are fewer than k nodes preceding μ, then return the nodes found. If μ is the head of its path then return *nil*.
- *after*(*node* μ, *integer* k) — return the k nodes succeeding μ on *getpath*(μ). If there are fewer than k nodes succeeding μ, then return the nodes found. If μ is the tail of its path then return *nil*.
- *precedes*(*node* μ, ν) — return **true** if nodes μ and ν are on the same path with μ preceding ν. Otherwise, returns **false**.
- *pathreport*(*path* Π) — report the value of path attribute set $\mathcal{P}$ for path Π.
- *nodereport*(*node* μ) — report the value of node attribute set $\mathcal{N}$ for node μ.
- *concatenate*(*path* Π', Π'') — concatenate paths Π' and Π'' to get path Π.
- *split*(*node* μ) — split path Π into three subpaths by removing the edges from μ to *before*($\mu, 1$) and μ to *after*($\mu, 1$).
- *reverse*(*path* Π) — reverse the direction of path Π, i.e., set $\Pi = \overline{\Pi}$.
- *exchange*(*path* Π) — exchange the values $\mathcal{N}(\mu)$ and $\mathcal{N}^*(\mu)$ for each node μ on Π.
- *localupdate*(*node* μ, *nodeattribute* N, *value* x) — set the value of node attribute N of $\mathcal{N}$ to be x.
- *globalupdate*(*path* Π, *nodeattribute* N, *value* δ) — for each node μ on path Π, set $N(\mu) = N(\mu) \odot \delta$, where N is a globally updatable node attribute with monoid operator $\odot$.
- *find*(*path* Π, *selectionfunction* S, *value* q) — find node μ of Π returned by the path-query associated with path-selection function S using query argument q.

Our data structure consists of representing each path Π as a balanced binary tree $\mathcal{B}_\Pi$, called the *path-tree* of Π, and identifying Π by the root ζ of $\mathcal{B}_\Pi$. Each leaf λ of $\mathcal{B}_\Pi$ represents a node of Π, and the left-to-right order of the leaves of $\mathcal{B}_\Pi$ corresponds to the head-to-tail order of the nodes of Π. We store at λ the values $\mathcal{N}(\lambda)$ and $\mathcal{N}^*(\lambda)$. For a globally updatable node attribute N_i of $\mathcal{N}$, the value stored at λ is not the actual value of N_i, but a "partial value" that can be used to compute $N_i(\lambda)$, as will become apparent later.

Each internal node η of $\mathcal{B}_\Pi$ represents the subpath $\Pi(\eta)$ of Π associated with the leaves in the subtree of η. Note that we identify path $\Pi(\eta)$ with η. We store at η pointers *head*(η) and *tail*(η) to the head and tail of $\Pi(\eta)$. Additionally, we store at node η the four values $\mathcal{P}(\Pi(\eta))$, $\mathcal{P}(\overline{\Pi(\eta)})$, $\mathcal{P}^*(\Pi(\eta))$, and $\mathcal{P}^*(\overline{\Pi(\eta)})$. Notice that we can consider the pointers *head* and *tail* as part of path attribute set $\mathcal{P}$. For the path attributes that depend on globally updatable node attributes, we store partial values. We also keep in η a vector $\Delta(\eta)$ (and a secondary vector $\Delta^*(\eta)$) such that for any globally updatable node attribute N_i of $\mathcal{N}$ the i-th component $\Delta_i(\eta)$ represents a change to the value of N_i for each node μ on $\Pi(\eta)$. (If N_j is a locally updatable node attribute, then we conventionally set $\Delta_j(\eta) = 0$.) Therefore, for any node μ of Π, $N_i(\mu)$ is the combination (using operation $\odot_i$, the monoid operator for values of N_i) of the "partial value" of N_i stored at μ and of the "offsets" $\Delta_i(\eta)$ of all the nodes η

on the path from μ to the root ζ of $\mathcal{B}_\Pi$. Similarly, $\mathcal{P}(\mu)$ can be computed from the partial value of $\mathcal{P}$ stored at μ and the "offsets" $\Delta(\eta)$ of all the nodes η on the path from μ to ζ in $\mathcal{B}_\Pi$. (Analogous arguments apply to the reverse and secondary values of the attributes.)

At each node μ of a path Π, we store pointers $left(\mu)$ and $right(\mu)$ to the left and right neighbors of μ on Π. At each path-tree node η we store pointers to $head(\Pi(\eta))$ and $tail(\Pi(\eta))$. We keep at node μ two globally updatable bits $reversed(\mu)$ and $exchanged(\mu)$. At each path tree node η, we keep the offset values $reversed(\eta)$ and $exchanged(\eta)$. If $reversed(\mu) = 1$, then $left(\mu)$ points to the right neighbor of μ and $right(\mu)$ points to the left neighbor of μ. Similarly, if $exchanged(\mu) = 1$, then the meaning of $\mathcal{N}(\mu)$ and $\mathcal{N}^*(\mu)$ is swapped. The meaning of the partial values $reversed(\eta)$ and $exchanged(\eta)$ is similar. If $reversed(\eta) = 1$, then $head(\eta)$ points to $tail(\Pi(\eta))$ and $tail(\eta)$ points to $head(\Pi(\eta))$. Also, $reversed(\eta) = 1$ swaps the meaning of the left and right child pointers in the subtree of η. The values of these two bits also determine the actual meaning of the four values kept for $\mathcal{P}$ at η.

Operations *makepath* and *deletepath* are trivially implemented in $O(1)$ time. Operation $pathreport(\Pi)$ is performed in $O(1)$ time and consists of returning the value of $\mathcal{P}$, computed from the "partial value" stored at the root ζ of $\mathcal{B}_\Pi$ and the vector $\Delta(\zeta)$. Operations $reverse(\zeta)$ and $exchange(\zeta)$ are implemented by flipping the *reversed* and *exchanged* bits at η. Operations $head(\zeta)$ and $tail(\zeta)$ determine the correct node based on the value of $reversed(\zeta)$ and the *head* and *tail* pointers stored at ζ. Operation $globalupdate(\Pi, N_i, \delta)$ is performed by setting $\Delta_i(\zeta) = \Delta_i(\zeta) \odot \delta$, where ζ is the root of the path-tree of Π. Each of these operations is implemented in $O(1)$ time.

Operation $getpath(\mu)$ is implemented by following pointers from μ to the root of the path-tree containing μ. Operation $before(\mu, k)$ is implemented by calling operation *push* on each node of the path from $getpath(\mu)$ to μ. The immediate predecessor of μ is then pointed at by $left(\mu)$. We use the *left* and *right* pointers in the following way to return the remaining (at most) $k-1$ predecessors of μ: Suppose we have returned $i < k$ nodes with μ_{i-1} and μ_i the $i-1$st and ith nodes returned. Then one of $left(\mu_i)$ and $right(\mu_i)$ will point to μ_{i-1}. The other will point to the predecessor of μ_i. (We cannot simply follow the *left* pointers, since each μ_i only stores a partial value of $reversed(\mu_i)$.) The implementation of operation $after(\mu, k)$ is symmetric. Each of these operations takes $O(d_\mu + k)$ time where d_μ is the depth of μ in $\mathcal{B}_\Pi$. Note that for a path of length ℓ, d_μ is $O(\log \ell)$.

Operation $precedes(\mu, \nu)$ is implemented by finding the least common ancestor η of μ and ν in the path tree containing both μ and ν. We then calculate the reversal state of η to determine which of μ and ν is a descendant of the left child of η. If μ is a descendant of the left child then return **true** else return **false**.

The following operation moves the values of Δ towards the leaves of a path-tree in $O(1)$ time:

- $push(node\ \eta)$ — For internal path-tree node η with children η' and η'' combine $\Delta(\eta')$ and $\Delta(\eta'')$ with $\Delta(\eta)$ and set $\Delta(\eta) = 0$.

Operation $push(\eta)$ is implemented by using the monoid operator for each globally updatable node attribute to apply the partial value at η to the partial values at η' and η''. We also swap values depending on $reversed(\eta)$ and $exchanged(\eta)$. If $reversed(\eta) = 1$, then we swap pointers $head(\eta)$ and $tail(\eta)$, the values $\mathcal{P}(\eta)$ and $\overline{\mathcal{P}(\eta)}$, and the values $\mathcal{P}^*(\eta)$ and $\overline{\mathcal{P}^*(\eta)}$. If $exchanged(\eta) = 1$, then we swap the values $\mathcal{P}(\eta)$ and $\mathcal{P}^*(\eta)$, and the values $\overline{\mathcal{P}(\eta)}$

and $\overline{\mathcal{P}^*(\eta)}$. If η is a tree node μ, then it is a leaf of its path tree. If $reversed(\mu) = 1$, we swap pointers $left(\mu)$ and $right(\mu)$. If $exchanged(\mu) = 1$, then we swap the values $\mathcal{N}(\mu)$ and $\mathcal{N}^*(\mu)$. (Equivalently, we can include the pointers *left* and *right* in node attribute set $\mathcal{N}$ and the pointers *head* and *tail* in path attribute set $\mathcal{P}$. In this way, operation *push* would only swap attribute sets, not individual pointers.)

Therefore, operation $nodereport(\mu)$ is implemented by finding the root ζ of the path-tree containing μ and then calling *push* along this path starting at ζ. In this way, the partial values of the globally updatable node attributes of $\mathcal{N}(\mu)$ will be the actual values. Therefore, we then return $\mathcal{N}(\mu)$.

Operations *concatenate* and *split* use the following *elementary tree operations* which can be performed in $O(1)$ time:

- $join(node\ \zeta', \zeta'')$ — Given the roots ζ' and ζ'' of the path-trees for solid paths Π' and Π'', combine the trees into a new tree by creating a new root ζ with left child ζ' and right child ζ''. We set the value of $\Delta(\zeta)$ to be 0. We set the *left* and *right* neighbor pointers at $tail(\zeta')$ and $head(\zeta'')$ to point at each other. Choose the one of *left* or *right* that is initially *nil* at $tail(\zeta')$ and set that value to $head(\zeta'')$. If both *left* and *right* are *nil* at $tail(\zeta')$ then we can determine in constant time the value of *reversed* at $tail(\zeta')$, so we set the *right* pointer to be $head(\zeta')$. We perform a similar update at $head(\zeta'')$. We calculate the values of $\mathcal{P}$ stored at ζ in $O(1)$ time using concatenation function F.
- $separate(node\ \zeta)$ — Given the root ζ of a binary tree, first call $push(\zeta)$ then divide the tree into two trees with roots ζ' and ζ'', where ζ' is the root of the left subtree and ζ'' is the root of the right subtree. The values of $\mathcal{P}$ at ζ' and ζ'' do not change.
- $rotateleft(node\ \eta)$ ($rotateright(node\ \eta)$) — Perform a left (right) rotation at node η. This operation can be performed with $O(1)$ *separate* and *join* operations.

We conclude operation $split(\mu)$ by calling $push(\mu)$. This step is not needed to maintain the values of paths, yet is needed in our implementation of trees (see section 2.2).

Operation $localupdate(\mu, x)$ is implemented by first replacing the value of $\mathcal{N}(\mu)$ with x. Suppose path tree $\mathcal{B}$, with root ζ, contains node μ. Suppose η is a path tree node in $\mathcal{B}$. The value of $\mathcal{P}(\eta)$ may change only if η is on the path from μ to ζ in $\mathcal{B}$. Therefore, we calculate these values starting at μ and ending at ζ. This takes $O(d_\mu)$ time where d_μ is the depth of μ in $\mathcal{B}$.

We need the following operation from [29] in our implementation of $find(\Pi, S, q)$:

- $splaystep(node\ \eta)$ — For binary tree $\mathcal{B}$ with root ζ and node η a grandchild of ζ, restructure $\mathcal{B}$ such that the relative order of the leaves of $\mathcal{B}$ remain fixed, and every node in the subtree rooted at η has its depth reduced by one or two. We accomplish this as follows. Let η' be the child of ζ that is the parent of η. If η' and η are both left children then perform $rotateright(\zeta)$ twice. If η' is a left child and η is a right child, then perform $rotateleft(\eta')$ followed by $rotateright(\zeta)$. The other two cases are symmetric. (See Fig. 1.)

Operation *splaystep* is implemented with a constant number of *rotateleft* and *rotateright* operations. Hence, *splaystep* takes $O(1)$ time.

Lemma 1 *Suppose S is a path-selection function for path attribute set $\mathcal{P}$ and ζ is the root of a path-tree Π. Then given query argument q, we can determine in $O(1)$ time if $S(\Pi, q)$*

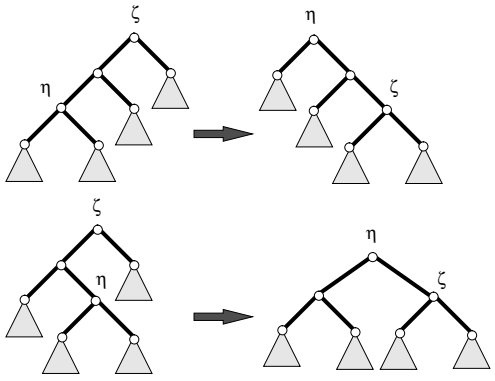

Figure 1: The rotations performed in operation *splaystep*(η).

is a child or grandchild of ζ*, or which grandchild of* ζ *is the root of the subtree containing* $S(\Pi, \zeta)$.

Proof: Suppose η_1 and η_2 are the children of ζ. By definition, in $O(1)$ time we can determine which subtree rooted at η_1 or η_2 contains $S(\Pi, q)$. If the subtree found is a single node, then it is $S(\Pi, q)$. Without loss of generality, suppose we know that $S(\Pi, q)$ is in the subtree rooted at η_1, and η_1 is not a leaf. Then let η_{11} and η_{12} be the children of η_1. We then issue *splaystep*(η_{11}) to restructure $\mathcal{B}_\Pi$ such that η_{11} is the left child of the root. Then, in $O(1)$ time we can determine if $S(\Pi, q)$ is in the subtree rooted at η_{11}. Otherwise, $S(\Pi, q)$ is in the subtree rooted at η_{12}. □

Operation *find*(ζ) is implemented as follows, starting at node ζ and repeating until a value is returned: if μ is a child or grandchild of ζ then return μ. Otherwise, determine the subtree rooted at the grandchild η of ζ which contains μ. Do *splaystep* at η and recur. After μ is found, we undo all the *splaystep*s to restore the balance of the path-tree. This takes $O(d_\mu)$ time where d_μ is the depth of the returned node μ.

Theorem 1 *Let* $\mathcal{Q} = (\mathcal{P}, \mathcal{N}, F)$ *be a path attribute system and* $\mathcal{S}$ *a collection of path-selection functions for* $\mathcal{Q}$*. There is a fully dynamic data structure for maintaining* $\mathcal{Q}$ *and* $\mathcal{S}$ *over a set of paths with the following performance:*

1. *a path of length* ℓ *uses* $O(\ell)$ *space;*
2. *operations* makepath, deletepath, head, tail, reverse *and* exchange *each take* $O(1)$ *time;*
3. *operations* before *and* after *each take* $O(\log \ell + k)$ *time to return* k *nodes from a path of length* ℓ*;*
4. *operations* getpath, concatenate, split, precedes, localupdate, globalupdate, pathreport, nodereport *and* find *each take* $O(\log \ell)$ *time for a path of length* ℓ*.*

Example 2 *Suppose we associate a value* cost(μ) *with each node* μ *of a path* Π*. We want to answer the following query in a dynamic environment (including operations* reverse, concatenate, *and* join*):*

- FindMin(*path* Π) — *return the minimum cost node of path* Π *closest to* head(Π)*.*

We keep the following value for each path Π*:*

- mincost(*path* Π) — *the minimum cost of any node of path* Π*.*

Therefore, we keep path attribute system $\mathcal{Q} = (\mathcal{N}, \mathcal{P}, F)$*, where for node* μ *and path* Π*,* $\mathcal{N}(\mu) = (\text{cost}(\mu))$ *and* $\mathcal{P}(\Pi) = (\text{mincost}(\Pi))$*. Suppose path* Π *is the concatenation of paths* Π' *and* Π''*. Concatenation function* F *calculates* $\text{mincost}(\Pi) = \min(\text{mincost}(\Pi'), \text{mincost}(\Pi''))$*.*

Operations FindMin *is the selection query associated with the following path selection function.*

Selection Function S_1 Suppose η is the root of a path tree with children η' and η''. If $mincost(\eta') = mincost(\eta)$ then return η', else return η''.

Operation FindMin(Π) *is then implemented as* find(Π, S_1)*.*

Example 3 *Suppose each node is associated one of a finite number of possible types. We wish to support the following query.*

- FindLongest(*path* Π, *nodetype* X) — *return the tail of the longest subpath* Π' *of* Π *such that* $\text{head}(\Pi') = \text{head}(\Pi)$ *and* Π' *does not contain a node of type* X*.*

We keep the following path attribute system $\mathcal{Q} = (\mathcal{N}, \mathcal{P}, F)$*. For each node* μ*, node attribute set* $\mathcal{N}(\mu)$ *consists of the single node attribute* $\text{type}(\mu)$*, the type of node* μ*. For path* Π*, path attribute set* $\mathcal{P}(\Pi)$ *consists of the set of boolean valued path attributes* $\text{contains}(\Pi, X)$*. The value of* $\text{contains}(\Pi, X)$ *is* **true** *if path* Π *contains a node of type* X*. Suppose path* Π *is the concatenation of paths* Π' *and* Π''*. Concatenation function* F *sets* $\text{contains}(\Pi, X)$ *to be* **true** *if either* $\text{contains}(\Pi', X)$ *is* **true** *or* $\text{contains}(\Pi'', X)$ *is* **true***.*

We utilize the following path selection function with query argument X*.*

Selection Function S_2 Suppose η is the root of a path tree with children η' and η''. If $contains(\eta', X)$ is **true**, or if $head(\eta'')$ has type X, then return η', else return η''.

Query FindLongest(Π, X) *is then implemented as follows. If* $\text{head}(\Pi)$ *has type* X*, then return nil. Otherwise, return* $\text{find}(\Pi, S_2, X)$*.*

Example 4 We consider a problem that arises in dynamically maintaining trees (see Section 2.2). Consider paths whose nodes have positive real weights. A node μ is *light* if the weight of μ is greater than the sum of the weights of the predecessors of μ on Π. Each path Π has at least one light node since $head(\Pi)$ is light. We want to answer the following queries in a dynamic environment (including operations *reverse*, *concatenate*, and *join*):

- $totalpathwt(path\ \Pi)$ — return the total weight of all the nodes in path Π.
- $light(path\ \Pi)$ — return the light node μ of Π closest to $tail(\Pi)$.

Let the *tilt* of a node μ of path Π be the sum of the weights of the nodes preceding μ in Π, minus the weight of μ. Therefore a node μ is light if $tilt(\mu) < 0$. Let $mintilt(\Pi)$ be the minimum value of *tilt* among all nodes of Π. Notice that $totalpathwt(\Pi) = totalpathwt(\overline{\Pi})$, but in general $negtiltnode(\Pi)$ is different from $negtiltnode(\overline{\Pi})$.

To solve this problem we keep the following path attribute system $\mathcal{Q} = (\mathcal{N}, \mathcal{P}, F)$: the node attribute of a node μ is the weight of μ, i.e., $\mathcal{N}(\mu) = (weight(\mu))$; the path attributes of a path Π are $\mathcal{P}(\Pi) = (totalpathwt(\Pi), mintilt(\Pi))$ (see Fig. 2). Suppose path Π is the concatenation of paths Π' and Π''. Concatenation function F calculates $\mathcal{P}(\Pi)$ as follows:

$$\begin{aligned} totalpathwt(\Pi) &= totalpathwt(\Pi') + totalpathwt(\Pi'') \\ mintilt(\Pi) &= \min(mintilt(\Pi'), totalpathwt(\Pi') - weight(head(\Pi'')), \\ &\quad totalpathwt(\Pi') + mintilt(\Pi'')) \end{aligned}$$

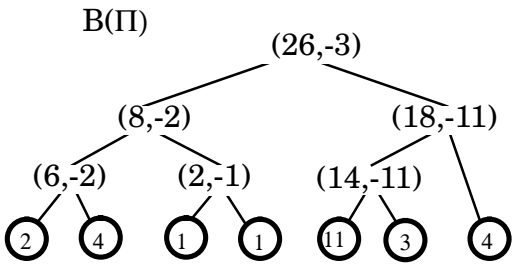

Figure 2: Path-tree for the path attribute system of Example 4. Each leaf is labeled with its weight. Each internal node η is labeled with $\mathcal{P}(\eta) = (totalpathwt(\eta), mintilt(\eta))$.

Operation $totalpathwt(\Pi)$ is implemented with $pathreport(\Pi)$. Operation $light(\Pi)$ is implemented as a selection-query via the following path-selection function without query argument.

Selection Function S_3 *Suppose η is the root of a path tree with children η' and η''. if* $\text{totalpathwt}(\Pi') - \text{weight}(\text{head}(\Pi'')) < 0$ *or* $\text{totalpathwt}(\Pi') + \text{mintilt}(\Pi'') < 0$ *then return η''. Otherwise, return η'.*

2.2 Tree Attribute Systems

We now discuss a general class of dynamic algorithms on trees. We consider rooted ordered trees with edges directed from the child to parent. Hence, a path in a tree is always directed from a descendant to an ancestor.

Definition 7 *Rooted tree T has* bounded degree *if there is a constant d such that no node of T has more than d neighbors, otherwise it is of* unbounded degree. *Tree T is* ordered *if for each node μ the left to right order of the children of μ is fixed.*

If T is a tree and μ is a node of T we use the notation T_μ to indicate the subtree rooted at μ.

We develop our dynamic algorithm for trees in the same manner as for paths. We introduce the concept of tree attribute systems, and show a large collection of dynamic operations on a tree storing a tree attribute system.

Definition 8 *A node attribute set $\mathcal{N}$ for a tree T consists of a finite collection of node attributes for the nodes of T. Similarly, a path attribute set $\mathcal{P}$ for T consists of a fixed collection of path attributes for the paths of T. A* tree attribute *R is a function on trees such that, for a tree T, $R(T)$ depends on the values $\mathcal{N}(\mu)$ for each node μ on T, on the parent-child relationships between nodes of T, and on the order of the children of each node of T. We assume that $R(T)$ can be stored in $O(1)$ space. A* tree attribute set *$\mathcal{R}$ is a fixed collection of tree attributes. For a node μ of tree T, $\mathcal{R}(\mu)$ denotes the tree attribute set of the subtree of T rooted at μ.*

We consider both bounded and unbounded degree trees. We first present the algorithm for bounded degree trees, then extend this result to unbounded degree trees.

2.2.1 Bounded Degree Trees

In this section, we present a fully dynamic algorithm for tree attribute systems on bounded degree trees. We begin by formally defining a tree attribute system.

Definition 9 *Suppose $\mathcal{N}$ is a node attribute set, $\mathcal{P}$ is a path attribute set, $\mathcal{R}$ is a tree attribute set, and F is a concatenation function. Given a bounded degree tree T and a path Π of T, we denote with $\tilde{\mathcal{N}}$ the extended node attribute set that, for a node μ of Π, consists of $\mathcal{N}(\mu)$, $\mathcal{R}(\mu')$ for each child μ' of μ not on Π, and the ordering of the children of μ. We say that the triplet $\mathcal{Z} = (\mathcal{N}, \mathcal{P}, \mathcal{R}, F)$ is a* tree attribute system *if:*

- *The triple $(\tilde{\mathcal{N}}, \mathcal{P}, F)$ is a path attribute system for the set of paths of T (where concatenations are restricted to paths Π' and Π'' in T such that* head(Π'') *is the parent of* tail(Π') *in T).*
- *The value of $\mathcal{R}$*(head(Π)) *can be determined in constant time from the values $\mathcal{P}(\Pi)$, $\tilde{\mathcal{N}}$*(head(Π))*, and $\tilde{\mathcal{N}}$*(tail(Π))*.*

We extend the concept of selection to trees. That is we define a tree selection function, which, if available, allows us to quickly find a distinguished node of tree T.

Definition 10 *Given a tree attribute system $\mathcal{Z} = (\mathcal{N}, \mathcal{P}, \mathcal{R}, F)$, a* tree-selection query *Q for $\mathcal{Z}$ maps tree T and a query argument q into a node $\mu = Q(T, q)$ of T. Suppose path Π is the concatenation of Π' and Π'' in T such that* tail(Π) *is the root of T. Suppose node μ' is the node closest to* tail(Π) *such that μ is a descendant of μ'. A* tree-selection function *$S(\Pi, q)$ for Q is a path-selection function to find node μ'.*

Definition 11 *Consider a tree T containing a node ν with children $\mu_1, \ldots, \mu_d$ in left-to-right order. Suppose Π is a path of T containing both ν and some child μ_i. The* left-tree *of ν with respect to Π $T_\ell(\nu)$ is defined as follows. The root of $T_\ell(\nu)$ is a node ν_ℓ with $\mathcal{N}(\nu_\ell) = \mathcal{N}(\nu)$. Root ν_r has children $\mu_1, \ldots, \mu_{i-1}$. The* right-tree *of ν with respect to Π $T_r(\nu)$ is defined similarly. The root of $T_r(\nu)$ is a node ν_r with $\mathcal{N}(\nu_r) = \mathcal{N}(\nu)$. Root ν_r has children $\mu_{i+1}, \ldots, \mu_d$.*

We provide three types of global restructuring of trees: reflecting, everting and cycling. Consider a tree T. We allow a node μ of T to be *fixed* indicating that the structure of the subtree rooted at μ may not be restructured.

Definition 12 *Given a tree T and a node μ of T that is not fixed, reflecting T at μ consists of reversing the left-to-right order of the children of μ and the children of all descendants that can be reached from μ without passing through a fixed node.*

Definition 13 *Given a tree T with node μ not the root,* everting *T at μ consists of reversing the parent-child relationships between all nodes on the path Π from μ to the root of T. If Π contains a fixed node, then everting at μ is not allowed.*

Each node μ also keeps an eversion rule *which describes changes to the order of the children of μ. Possibilities include:*

- *Simple — the left and right-trees of μ remain unchanged.*
- *Steady — the left and right-trees of μ are exchanged.*
- *Mirrored — the left and right-trees of μ are reflected.*

Consider a node μ, its child μ', and its parent ν, all on path Π. Simple eversion at μ substitutes ν for μ' in the ordering of the children of μ. Steady eversion at μ maintains the clockwise order of the neighbors of μ. Finally, mirrored eversion at μ reverses the clockwise order of the neighbors of μ. Each node of a tree T may use a different eversion rule. Therefore, we consider the eversion rule a node property.

Definition 14 *Given a tree T and a node μ and its parent ν of T that is not fixed, cycling ν at μ consists of cyclicly permuting the children of ν such that node μ is the final node in the order.*

Our algorithm supports the following operations on a collection of trees with a tree attribute system $\mathcal{Z} = (\mathcal{N}, \mathcal{P}, \mathcal{R})$ and a collection $\mathcal{F}$ of path-selection and tree-selection functions:

- *Evaluate*(*node* μ) — Return the values of $\mathcal{N}(\mu)$ and $\mathcal{R}(\mu)$.
- *PathFind*(*node* μ', μ''; *selectionfunction* S) — Let Π be the path of a tree from node μ' to node μ''. Find the node of Π returned by the path-selection function S.
- *TreeFind*(*node* ν; *selectionfunction* S) — Find the node of the subtree rooted at μ returned by the tree-selection function S.
- *LocalUpdate*(*node* μ, *nodeattribute* N, *value* x) — Set the value of node attribute N of $\mathcal{N}$ to be x.
- *GlobalUpdate*(*node* μ', μ'', *nodeattribute* N, *value* δ) — Let Π be the path from μ' to μ'' in T. For each node μ on path Π, set $N(\mu) = N(\mu) \odot \delta$, where $\odot$ is the operator of the monoid for N.
- *MakeTree*(*value* x) — Create a new tree consisting of a single leaf whose node attribute set has value x.
- *DeleteTree*(*node* λ) — Remove the single-node tree consisting of leaf λ.
- *LCA*(*node* μ', μ'') — Returns the least common ancestor of nodes μ' and μ''. Nodes μ' and μ'' are assumed to be nodes of the same tree.
- *Evert*(*node* μ) — Let T be the tree containing node μ. Evert T at μ.
- *Link*(*node* ρ, μ, μ', μ'') — This operation assumes that nodes μ and ρ are in different trees, and ρ is the root of its tree T. Add an edge from ρ to μ with μ' its immediate left sibling and μ'' its immediate right sibling. If μ' or μ'' is *nil* then ρ becomes the first or last child of μ. Tree T becomes a subtree of the tree containing μ.
- *Cut*(*node* μ) — This operation assumes that μ is not the root of a tree. Remove the edge from μ to its parent, thus separating the subtree rooted at μ.
- *Parent*(*node* μ) — Return the parent of μ or *nil* if μ is a tree root.
- *Sibling*(*node* μ, *integer* k, *direction* X) — Return the k adjacent siblings of node μ in the X direction, where X is either *left* or *right*. If there are fewer than k siblings in the X direction, then return the nodes found. If there are no siblings in the X direction, then return *nil*.
- *Reflect*(*node* μ) — Reflect the subtree rooted at μ.
- *Cycle*(*node* ν, μ) — Cyclicly permute the children of ν such that node μ is the final node in the order. Node μ is assumed to be a child of node μ.

η'	η''	$\eta' \odot_r \eta''$
$(x,0)$	(y,z)	$(x+y,0)$
$(0,1)$	(y,z)	$(0,1)$

Table 1 The rules to calculate the values of operator $\odot_r$. Variables x, y, and z take boolean values. Operator + is boolean addition.

Our data structure consists of representing an n-node tree T as a collection of disjoint paths. We partition the edges of T into *solid* or *dashed* such that at most one solid edge is incoming into a node. Therefore, every node is in exactly one maximal path of solid edges (of length 0 or more), called a *solid path* of T.

The partition of edges into solid and and dashed is obtained by means of the following *size invariant*: let $size(\mu)$ denote the number of nodes in the subtree of T rooted at node μ. An edge (μ,ν) of T is solid if $size(\mu) > size(\nu)/2$. If the partition satisfies the size invariant, then every path of T has $O(\log n)$ dashed edges.

Also, in order to achieve the logarithmic time per dynamic operation, we use biased search trees [2] to represent the path-trees. For a node ν of a solid path Π we define the node-attribute $weight(\nu)$ as one plus the sum of the sizes of the subtrees connected to μ by dashed edges. For each solid path Π, we also define the corresponding path attribute $weight(\Pi)$ as the sum of the weights of the nodes of Π. The path-tree $\mathcal{B}_\Pi$ is then implemented as a search tree biased by the node-weights of its leaves.

The size invariant may be temporarily violated by the algorithms that operate on the solid paths of T. Hence, we set-up a mechanism to restore it, using a path attribute system on the solid paths.

Consider a node μ of T on solid path Π. We keep a globally updatable node attribute $reflected(\mu)$ which indicates the *path reflection status*, which is the left-to-right direction of the children of μ, ignoring any reflection done by ancestors of $tail(\Pi)$. At each path tree node η, we keep the associated offset value $reflected(\eta)$. Rather than use the *exchange* bits for paths, we use the *reflected* bits which are handled somewhat differently. We also keep a locally updatable node attribute $fixed(\mu)$ which indicates if node μ is fixed.

Even though the values for *reflected* are boolean, we do not use simple boolean addition as the monoid operator for calculating the actual value of $reflected(\mu)$, for a node μ. At a path tree node η, we consider $reflected(\eta)$ to be part of a pair of values $(a(\eta),\ b(\eta))$, where $a(\eta) = reflected(\eta)$ and $b(\eta) = fixed(tail(\eta))$. The monoid operator $\odot_r$ on these pairs is described in table 1. The value $(1,1)$ is not allowed since we do not allow reflection of a fixed node. Note that operator $\odot_r$ does not modify the value of $fixed(tail(\eta))$.

Let μ be a node of tree T. Let ν be the node closest to the root of T such that ν is reached by a dashed edge on the path Π from μ to the root of T, and ν is a descendant of the closest ancestor of μ that is fixed. If μ has no fixed ancestors, then μ will be the parent of the last dashed edge of Π. The *reflection status* of node μ is then calculated by finding the exclusive-or of the path reflection status of μ and the nodes reached over dashed edges on the path from μ to ν.

We use the *reflected* bits much like the *exchanged* bits of section 2.1. Let $\mathcal{Z} = (\mathcal{N},\mathcal{P},\mathcal{R})$ be the tree attribute system we are maintaining for tree T. For node μ, let $\mathcal{R}^*(\mu)$ be the tree

	Initial Value		*Restructured*	
Rule	*reversed*	*reflected*	$T_\ell(\mu)$	$T_r(\mu)$
any	0	0	$T_\ell(\mu)$	$T_r(\mu)$
any	0	1	$T_r^*(\mu)$	$T_\ell^*(\mu)$
simple	1	0	$T_\ell(\mu)$	$T_r(\mu)$
simple	1	1	$T_r^*(\mu)$	$T_\ell^*(\mu)$
steady	1	0	$T_r(\mu)$	$T_\ell(\mu)$
steady	1	1	$T_\ell^*(\mu)$	$T_r^*(\mu)$
mirrored	1	0	$T_\ell^*(\mu)$	$T_r^*(\mu)$
mirrored	1	1	$T_r(\mu)$	$T_\ell(\mu)$

Table 2 Restructuring done by operation $push(\mu)$ for tree node μ with left-tree $T_\ell(\mu)$ and right-tree $T_r(\mu)$. Eversion rules are given in definition 13. The notation $T_\ell^*(\mu)$ and $T_r^*(\mu)$ indicate the reflections of $T_\ell(\mu)$ and $T_r(\mu)$.

attribute set for the reflection of μ. We have defined the (primary) extended node attribute set $\tilde{\mathcal{N}}(\mu)$ to be the union of $\mathcal{N}(\mu)$ and $\mathcal{R}(\mu')$ for each child μ' of μ not on Π. The secondary extended node attribute set $\tilde{\mathcal{N}}^*(\mu)$ is the union of $\mathcal{N}(\mu)$ and $\mathcal{R}^*(\mu')$ for each child μ' of μ not on Π. Then for each path tree node η we keep the four values $\mathcal{P}(\Pi(\eta))$, $\mathcal{P}(\overline{\Pi(\eta)})$, $\mathcal{P}^*(\Pi(\eta))$, and $\mathcal{P}^*(\overline{\Pi(\eta)})$, the meaning of which depend on the values of $reversed(\eta)$ and $reflected(\eta)$.

We keep a value in $\mathcal{P}(\Pi)$ called $containsfixed(\Pi)$ indicating if path Π contains a fixed node. If so, then the values of $\mathcal{P}(\overline{\Pi(\eta)})$ and $\mathcal{P}^*(\overline{\Pi(\eta)})$ are undefined.

For node μ on path Π in tree T, we extend the restructuring done by operation $push(\mu)$. If the value of $reversed(\mu) = 1$ or $reflected(\mu) = 1$ prior to the call to $push(\mu)$, then we may swap the left and right-trees of μ, reflect the left and right trees of μ, or both. The restructuring done is based on the eversion rule for μ as described in definition 13. Table 2 describes this restructuring. For bounded degree tree T, we swap the left and right-trees of μ by reordering the children of μ. We reflect the left (right) tree of μ by reversing the order of the children of μ in the left (right) tree and then flipping the *reflected* bit at the roots of the path trees containing the children of μ in the left (right) tree. Since T has bounded degree, we can perform all this in $O(1)$ time.

Now, we show how to perform operation $Evaluate(\mu)$. Let Π be the solid path of length ℓ containing μ. If μ is the tail of Π then, by the definition of tree attribute systems, we can use path attribute system $(\tilde{\mathcal{N}}, \mathcal{P})$ to find $\mathcal{R}(\mu)$ in $O(1)$ time. Otherwise, we make μ the tail of a path by calling $split(\mu)$. We then restore path Π with the *concatenate* operation. Hence, this can all be accomplished in $O(\log \ell)$ time.

Operation $LocalUpdate(\lambda, x)$ affects only the values of the nodes in the path from λ to the root ρ. If such path contains dashed edges, then we temporarily convert it into a solid path by changing dashed edges to solid and solid edges to dashed such that any node continues to have at most one incoming solid edge. This is done with the following operations, derived from dynamic trees [28]:

- $splice(path\ \Pi)$ — This operation assumes that Π is a solid path ending at $\mu \neq \rho$. Convert the dashed edge leaving μ to solid and convert the solid edge (if it exists)

entering the parent ν of μ to dashed. Let Π' be the solid path containing ν and μ' be the sibling of μ with the solid edge (μ', ν). To implement *splice*(Π) we need to convert edge (μ', ν) from solid to dashed. Let Π'' be the resulting solid path starting at ν. We obtain Π'' by splitting Π'. We complete *splice*(Π) by concatenating Π and Π''.

- *expose*(μ) — Convert to dashed the solid edge entering μ, if such edge exists. Create a solid path from μ to the root by converting to solid all the dashed edges (ν', ν'') of such path, and converting to dashed the edges $(\text{sib}(\nu'), \nu'')$. Operation *expose*$(\mu)$ consists of a sequence of *splice* operations on the solid paths containing the nodes on the path from μ to ρ. This operation is always followed by a *conceal* operation, which undoes its effect.
- *conceal*(μ) — Restore the original type (solid or dashed) of the edges entering the nodes on the path Π from node μ to the root ρ. This operation is the inverse of *expose*, and also consists of a sequence of *splice* operations. The topmost node where to *splice* is given by *light*(Π), as described in Example 4. Operation *light*(Π) returns the node ν of Π closest to *tail*(Π) such that the weight of ν is larger than the sum of the weights of the nodes preceding ν in Π. As shown in Example 4, *light*(Π) is found using a path-selection function.

Since path-trees are biased by the weights, the sum of the number of elementary tree operations at each *splice* operation telescopes, so that operations *expose* and *conceal* can be performed with $O(\log n)$ elementary tree operations.

Operation *LocalUpdate*(μ, x) is realized as follows. First, we issue *expose*(μ). Next, we change the value of the constant node attribute at μ to x. Finally, we perform *conceal*(μ) to restore the original solid and dashed edges. Hence, operation *LocalUpdate* is implemented in $O(\log n)$ time.

Operation *GlobalUpdate*(μ', μ'', N, δ) first makes two calls to *expose* to get the path Π from μ' to μ''. Let Δ_N be the entry for node attribute N in the Δ vector at the root of the path-tree for Π. We set Δ_N to $\Delta_N \odot \delta$. We restore the weight invariant by calling *expose*(μ') followed by *conceal*(μ'). All this can be done in $O(\log n)$ time.

To implement $LCA(\mu', \mu'')$ we first issue *expose*(μ') followed by *expose*(μ''). If μ' and μ'' are on the same resulting solid path, then μ' is an ancestor of μ'', so we return μ''. Otherwise, when we *expose*(μ''), we split the path containing μ' at the least common ancestor. Therefore, $LCA(\mu', \mu'')$ will be the parent of the tail of the path containing μ'. We restore the size invariant by calling *conceal*(μ''), *expose*(μ'), then *conceal*(μ'). All this can be performed in $O(\log n)$ time.

For a path Π, we use the eversion rules described in table 2 to calculate the value of $P(\bar{\Pi})$ and $P^{(}\bar{\Pi})$. We perform operation *Evert*(μ) by first issuing *expose*(μ) to get the path Π from μ to the root ρ, flipping the reverse bit of Π, and finally calling operation *conceal*(ρ) to restore the size invariant. Therefore, operation *Evert* is performed in $O(\log n)$ time.

Operation *Reflect*(μ) is performed as follows. First, we call *expose*(μ). The left edge-path of μ then contains nodes for all the children of μ. We then reverse the order of the children of μ by flipping the *reflected* and *reversed* bits at the root of the path-tree representing $\Pi_\ell(\mu)$. Finally, we perform *conceal*(μ) to restore the original solid and dashed edges. Hence, operation *Reflect* is implemented in $O(\log n)$ time.

Operation $Cycle(\nu, \mu)$ is performed by first calling $expose(\nu)$. We then reorder the children of ν. Suppose $T_r(\nu)$ is the right-tree of ν with respect to path Π before the *Cycle* operation is performed. We flip the *Reflect* bit at the roots of the path trees containing the children of ν in $T_r(\nu)$. Finally, we perform $conceal(\nu)$ to restore the original solid and dashed edges. All this can be done in $O(\log n)$ time.

Operation $PathFind(\mu', \mu'', S, q)$ makes two calls to *expose* to get the path Π from μ' to μ'', then calls $find(\Pi, S, q)$. We then restore the weight invariant by calling $expose(\nu')$ followed by $conceal(\nu')$.

Consider tree selection function S. We implement operation $TreeFind(\nu, S, q)$ as follows. If ν is not the root of its tree, we first call *expose* at $Parent(\nu)$. Node ν then becomes the tail of its solid path Π. Let μ be the node we are searching for. We repeat the following until node μ is found. Let node $\mu' = find(\Pi, S_p, q)$. Create path Π' with head μ' and tail ν. This is done with single *split* and *concatenate* operations.

Suppose $\mu_1, \ldots, \mu_d$ are the children of μ' not on Π and $\Pi_1, \ldots, \Pi_d$ are the solid paths containing each μ_j. Repeat the following until we find a child μ_i of μ' such that node μ is a descendant of μ_i. Consider child node μ_j. Let path Π' be the result of joining paths Π_j and Π'. By the definition of tree and path selection functions, we can determine in $O(1)$ time if node μ is in the subtree rooted at μ_j.

If no child of μ' has μ as a descendant, then $\mu' = \mu$. Otherwise, we reset path Π to be Π' and continue. After μ is found, we undo the restructuring to restore the path trees.

At each iteration, the time to find μ' and create Π' is $O(d_{\mu'})$, where $d_{\mu'}$ is the depth of μ' in $\mathcal{B}_\Pi$, the path tree for path Π. Operation *join* increases the depth of a path tree node by 1. Therefore, $d_{\mu'}$ is at most one more than the depth of μ' in its original path-tree. Therefore, the time to find μ is $O(\log n)$ plus the time to perform $expose(\mu)$. The time to restore the path trees is equivalent. Therefore operation *TreeFind* is implemented in $O(\log n)$ time.

Operations *MakeTree* and *DeleteTree* can be trivially implemented in $O(1)$ time. Operation $Parent(\mu)$ is performed in $O(\log n)$ time by calling $expose(\mu)$, returning $after(\mu)$ on the solid path containing μ, and calling $conceal(\mu)$.

Operation $Sibling(\mu, k, X)$ is realized as follows. If node μ is the root of T, then return *nil*. Otherwise, we call $expose(\mu)$. Let node ν be the parent of μ. Let ζ be the root of the resulting path tree. We call operation *push* on the path from ζ to ν. Node ν then stores the actual value of $reflected(\nu)$. We then return the (at most) k siblings of μ in the X direction. Hence, operation *Sibling* is performed in $O(\log n + k)$ time.

Operations *Link* and *Cut* are implemented in $O(\log n)$ time using variations of the same named dynamic tree operations. We implement operation $Link(\rho, \mu, \mu', \mu'')$ by first calling $expose(\mu)$. We then add ρ as a child of μ between μ' and μ''. Finally, we call $conceal(\mu)$ to maintain the size invariant.

Operation $Cut(\mu)$ is the inverse of *Link*. We first call $expose(Parent(\mu))$. We then remove node μ as a child of its parent by removing μ from the order of children stored at μ and destroying the dashed edge. We conclude by calling $conceal(\mu)$ to maintain the size invariant.

Theorem 2 *Let $\mathcal{Z} = (\mathcal{N}, \mathcal{P}, \mathcal{R}, F)$ be a tree attribute system, and $\mathcal{S}$ be a collection of path-selection and tree-selection functions. There is a fully dynamic data structure for maintaining $\mathcal{Z}$ and $\mathcal{S}$ over a set of trees with the following performance:*

1. *a tree of size n uses $O(n)$ space;*
2. *operations* MakeTree *and* DeleteTree *take each $O(1)$ time;*
3. *operations* Evaluate, LocalUpdate, GlobalUpdate, Link, Cut, Evert, Reflect, Cycle, *LCA,* PathFind, TreeFind, Sibling, *and* Parent *take each $O(\log n)$ time, where n is the total size of the trees involved in the operation.*

Example 5 *A* concatenable heap *consists of a collection of* heaps *which are sets of weighted nodes that support the following operations:*

- FindMin(*heap H*) — *returns the cost of a minimum cost node in heap H.*
- InsertHeap(*node μ; heap H*) — *insert node μ in heap H.*
- DeleteHeap(*node μ; heap H*) — *delete node μ from heap H.*
- ConcatenateHeap(*heap H_1, H_2*) — *combine heaps H_1 and H_2 into a single heap.*

We can implement a concatenable heap as either a path or a tree. If we store each heap as a path, then operation FindMin(H) *is implemented as operation* FindMin *in example 2. Operations* InsertHeap(μ, H), DeleteHeap(μ, H), *and* ConcatenateHeap(H_1, H_2) *are directly implemented using operations* concatenate *and* split.

The implementation of a concatenable heap as a tree is more complicated. We present this technique since it is used in developing more generalized heaps that can not be implemented as paths.

The nodes of a heap are stored as leaves of an arbitrary tree. A heap is identified as its representative tree. We maintain the following tree attribute for the nodes of tree T.

- mincost(μ) — *the minimum cost of a node in the subtree rooted at μ.*

We also maintain the following path attribute for the paths of tree T.

- minpathcost(Π) — *the minimum cost of a leaf that is a descendant of* tail(Π), *but is not* head(Π) *or a descendant of* head(Π).

It is easy to see that these values can be maintained as part of a tree attribute system. We implement operation FindMin *using the following tree selection function which takes a cost q as a query argument:*

Selection Function S_4 Suppose node η the root of the path tree with children η' and η''. If $minpathcost(tail(\eta)) = minpathcost(tail(\eta'))$ then return η', else return η''.

For tree T associated with heap H, let c be the value of mincost *at the root of T. Operation* FindMin(H) *is implemented as* TreeFind(ρ, S_1, c). *Operations* InsertHeap(μ, H), DeleteHeap(μ, H), *and* ConcatenateHeap(H_1, H_2) *are directly implemented using operations* Link, Cut, *and* MakeTree.

Therefore, both implementations of concatenable heaps are implemented with $O(\log n)$ time per operation, where n is the total size of the heaps involved.

2.2.2 Unbounded Degree Trees

The algorithm presented in the previous section for bounded degree trees takes advantage of the fact that reordering and reflecting the children of a node can be accomplished in constant time. In the case of unbounded degree trees, we need to develop additional techniques to accomplish the needed restructuring.

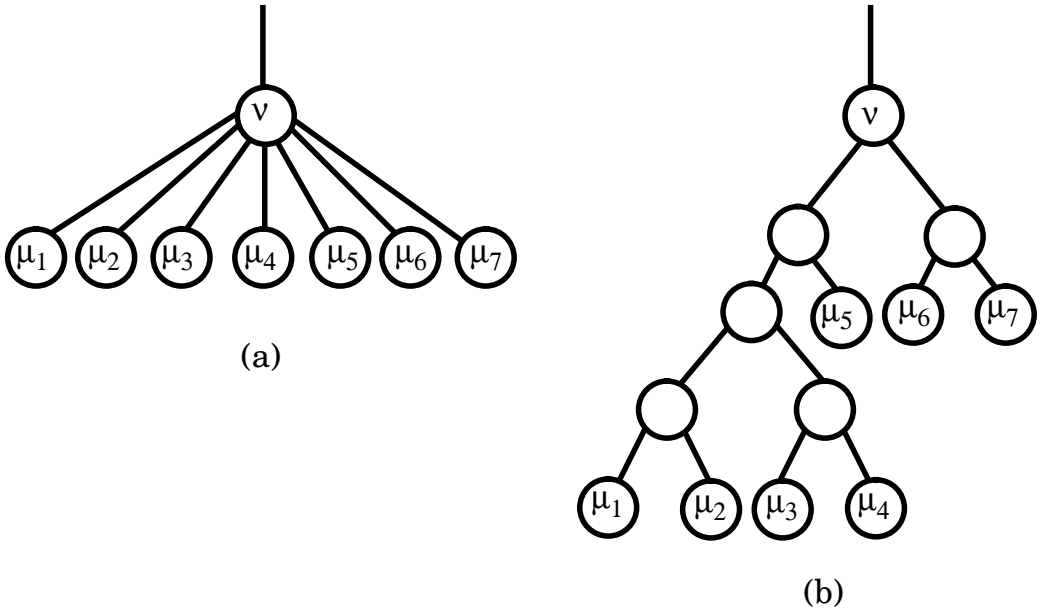

Figure 3: Expanding a node of unbounded degree. (a) The original node. (b) The expansion of the node.

Definition 15 *Suppose T is a (unbounded-degree) tree containing node ν with children $\mu_1, \ldots, \mu_d$.* Expanding *ν between μ_i and μ_j ($1 \leq i \leq j \leq d$) consists of restructuring T such that we replace nodes $\mu_i, \ldots \mu_j$ in the ordering of the children of ν with a new node μ. Node μ has children $\mu_i, \ldots \mu_j$. We set $\mathcal{N}_L(\mu) = \mathcal{N}_L(\nu)$ and $\mathcal{N}_G(\mu) = 0$. The inverse operation is called* contracting *child μ of ν. In order to contract μ, $\mathcal{N}_L(\mu)$ must equal $\mathcal{N}_L(\nu)$ and $\mathcal{N}_G(\mu)$ must equal 0.*

For unbounded degree trees, we require that tree attribute set $\mathcal{R}$ meet the following *expansion invariant*: For any node ν of tree T, the value of $\mathcal{R}(\nu)$ is unchanged as a result of expanding ν or contracting a child of ν. Therefore, given an unbounded degree tree T we can create an equivalent tree T' of bounded degree 2, called a *binary expansion* of T, by expanding nodes with more than two children (see Fig. 3). Note that if T has n nodes, then T' will have $O(n)$ nodes.

Definition 16 *Suppose T is an unbounded degree tree. Then $\mathcal{Z} = (\mathcal{N}, \mathcal{P}, \mathcal{R}, F)$ is a tree attribute system on T if and only if node attribute set $\mathcal{R}$ keeps the expansion invariant and $\mathcal{Z}$ is a tree attribute system on any binary expansion T' of T.*

For each node ν in an unbounded degree tree T, we keep two paths $\Pi_\ell(\nu)$ and $\Pi_r(\nu)$ associated with the children to the left and right of a specified child μ_i of ν (see Fig. 4). We write $\nu(\mu_j)$ to represent the node in either $\Pi_\ell(\nu)$ or $\Pi_r(\nu)$ associated with μ_j. Formally,

Definition 17 *Consider an unbounded degree tree T containing a node ν with children $\mu_1, \ldots, \mu_d$ in left-to-right order. Suppose Π is a path of T containing ν. The* left edge-path *of ν with respect to Π, $\Pi_\ell(\nu)$, consists of a node $\nu(\mu_j)$ for each child μ_j of ν in the left-tree of ν with respect to Π. Path $\Pi_\ell(\nu)$ is directed in the left-to-right order of the chil-*

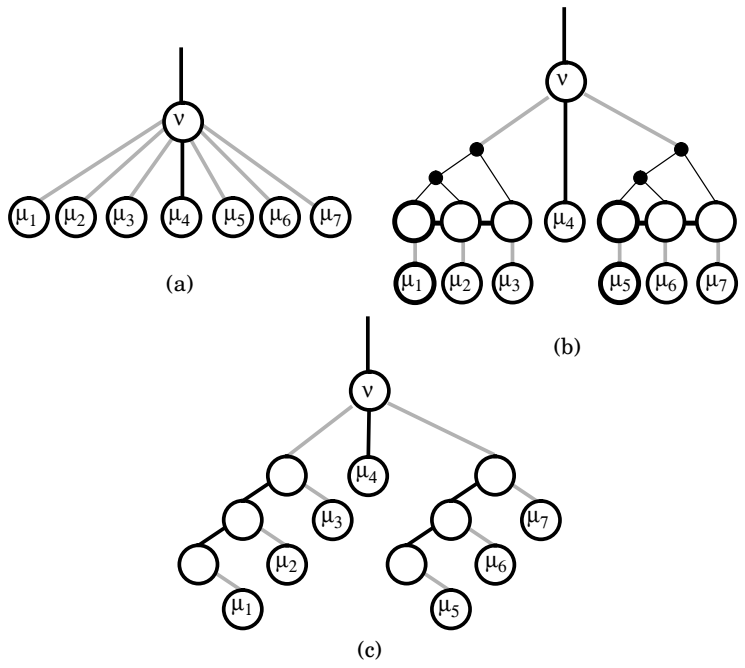

Figure 4: A node ν in an unbounded degree tree. (a) The node its solid path, and its children. (b) The left and right edge-paths. (c) An equivalent bounded degree representation.

dren. The value of $\mathcal{N}_L(\nu(\mu_j))$ *is defined to be* $\mathcal{N}_L(\nu)$. *The* right edge-path $\Pi_r(\nu)$ *is similarly defined for the children of* ν *in the right-tree of* ν.

Consider an unbounded degree tree T and tree attribute system $\mathcal{Z} = (\mathcal{N}, \mathcal{P}, \mathcal{R}, F)$ on T. Let Π be a solid path in T and $\Pi_X(\nu)$ be an edge path of some node ν. We define the extended node attribute set $\hat{\mathcal{N}}$ that, for a node μ of Π, $\hat{\mathcal{N}}(\mu)$ consists of $\mathcal{N}(\mu)$, $\mathcal{R}(\mu_\ell)$, and $\mathcal{R}(\mu_r)$ for the roots μ_ℓ and μ_r of the left and right-trees of μ. For a node $\nu(\mu')$ of $\Pi_X(\nu)$, $x = \ell$ or r, extended node attribute set $\hat{\mathcal{N}}(\nu(\mu'))$ consists of $\mathcal{N}(\nu)$ and $\mathcal{R}(\mu')$.

Lemma 2 *Suppose* $\mathcal{Z} = (\mathcal{N}, \mathcal{P}, \mathcal{R}, F)$ *is a tree attribute system on unbounded degree tree* T. *Then,*

- *The triplet* $(\hat{\mathcal{N}}, \mathcal{P}, F)$ *is a path attribute system for paths in* T *where concatenations are restricted to adjacent subpaths.*
- *The triplet* $(\hat{\mathcal{N}}, \mathcal{P}, F)$ *is a path attribute system for edge-paths where concatenations are restricted to adjacent subpaths of the edge-paths.*

Proof: Suppose Π is a solid path from node σ to node τ. Consider the path $\hat{\Pi}$ formed by expanding every node μ on Π into two nodes μ' and μ'' on $\hat{\Pi}$ with μ' the parent of μ''. Recall that nodes μ_ℓ and μ_r are the roots of the left and right-trees of node μ with respect to Π. Suppose we know the values $\mathcal{R}(\mu_\ell)$ and $\mathcal{R}(\mu_r)$. We make μ_ℓ the left child of μ' and μ_r the right child of μ''. Path $\hat{(\Pi)}$ is the path from σ'' to τ' in a binary expansion of T. Therefore, the first part of the lemma is true if we define $\mathcal{P}(\Pi)$ to be $\mathcal{P}(\hat{\Pi})$.

The argument for edge-paths is even more direct. Suppose ν is a node and X is either ℓ or r. Each node of edge-path $\Pi_X(\nu)$ has degree bounded by 2. Therefore each edge path is a path in a binary expansion of T, so the lemma holds by definition 16.

□

In addition to the operations previously defined for bounded degree trees, we support the following operations on unbounded degree trees:

- *Expand*(*node* μ, μ', μ'') — Expand node μ between children μ' and μ''. Node μ' is assumed to precede μ'' in the left-to-right order of the children of μ.
- *Contract*(*node* μ) — Contract μ into its parent node ν. The values $\mathcal{N}(\mu)$ and $\mathcal{N}(\nu)$ are assumed to be equal.

Edge-paths are stored in the same structure as solid paths. If $\Pi_X(\nu)$ is the left or right edge-path of node ν, then reversing $\Pi_X(\nu)$ is equivalent to reflecting the left or right-tree of ν. Therefore, path tree nodes use eversion rule 2 (see definition 13).

Recall that operation *splice* converts a dashed edge to solid and possibly a solid edge to dashed. Suppose μ is a node connected to its parent ν by a solid edge. Let Π be the solid path containing nodes μ and ν. Let $\Pi_\ell(\nu)$ and $\Pi_r(\nu)$ be the left and right edge-paths of ν. We convert solid edge (μ, ν) to dashed as follows. We begin by splitting path Π at ν. As described in section 2.2.1, operation *split* calls operation push along the path from the root of the path tree for Π to node ν. Operation $push(\nu)$ may require the left and right-trees of ν to be swapped or exchanged (see table 2). This is performed by either exchanging the meaning of $\Pi_\ell(\nu)$ and $\Pi_r(\nu)$, or by flipping the *reversed* and *reflected* bits at the root of the path tree representing $\Pi_\ell(\nu)$ and $\Pi_r(\nu)$. We conclude by creating a path tree node $\nu(\mu)$ for the new dashed edge, and setting the left edge-path of ν to be the concatenation $\Pi_\ell(\nu)$, $\nu(\mu)$, and $\Pi_r(\nu)$ and the right edge-path to be the empty path.

Converting a dashed edge to solid is performed similarly, reversing the roles of the solid and edge-paths. Suppose the edge between node μ and its parent ν is now dashed, and there is no solid edge entering ν. Let Π' and Π'' be the paths with $\mu = tail(\Pi')$ and $\nu = head(\Pi'')$. Let $\Pi_\ell(\nu)$ be the left edge-path of ν. Split $\Pi_\ell(\nu)$ at $\nu(\mu)$. The resulting paths to the left and right of $\nu(\mu)$ become the left and right edge-paths of ν. We then concatenate Π' and Π'' to create the solid edge.

For a node μ of a solid path Π we now define the node-attribute $weight(\mu)$ as one plus the sum of the sizes of the left and right edge-paths of μ. For a node $\nu(\mu')$ of an edge-path $\Pi_X(\nu)$ we define the node-attribute $weight(\nu(\mu'))$ as one plus the sum of the sizes of solid path containing μ'. We bias by weight the path trees for both the edge-paths and solid paths. Therefore, the sum of the number of elementary tree operations at each *splice* operation telescopes, so that operations *expose* and *conceal* can be performed with $O(\log n)$ elementary tree operations.

The techniques to perform operations *TreeFind*, *Sibling*, *Reflect* and *Cycle* are slightly modified for unbounded degree trees. All other operations remain unchanged.

We implement operation $TreeFind(\nu, S, q)$ as in the bounded-degree tree case, with following difference. Let μ be the node we a searching for. Suppose at an intermediate step node μ' is the deepest node on path Π such that the node μ is a descendant of μ'. We proceed as in the bounded case, except we next consider the edge-paths connected to μ', rather than the solid paths containing the children of μ'. Hence, we join an edge-path to Π. This is well defined by the expansion invariant. We perform a similar operation when μ' is an edge-path node.

Operation $Sibling(\mu, k, X)$ is realized by first calling $expose(\mu)$. We then use operation *push* to get the value of *reverse* and *Reflect* at $Parent(\mu)$. We find the siblings by either calling *before* or *after* on the left or right path trees of ν, depending on direction X. We conclude by restoring the size invariant by calling $conceal(\mu)$.

Operation $Reflect(\mu)$ is performed as follows. First, we call $expose(\mu)$. We then flip the *reverse* bit for the left edge-path of μ. Finally, we perform $conceal(\mu)$ to restore the original solid and dashed edges.

Operation $Cycle(\nu, \mu)$ is performed by first calling $expose(\nu)$. We then split the left path tree of ν at path tree node $\nu(\mu)$. Let Π' and Π'' be the resulting subpaths representing the children of ν before and after μ. We reorder the children of ν by flipping the *reflected* and *reversed* bits at the root of the path tree representing Π'', concatenating Π'', Π' and $\nu(\mu)$, and making this new edge-path the left edge-path of ν. Finally, we perform $conceal(\nu)$ to restore the original solid and dashed edges.

Operation $Expand(\mu, \mu', \mu'')$ begins by calling $expose(\mu)$ which results in the left edge-path of μ, $\Pi_\ell(\mu)$ containing all the edges from children to μ. Then we perform a constant number of split and join operations to form three subpaths Π', Π'', and Π''' of $\Pi_\ell(\mu)$, with Π' containing the edge-path nodes associated with the children of μ preceding μ, Π'' containing the edge-path nodes associated with the children of μ between μ' and μ'', and Π''' containing the edge-path nodes associated with the children of μ following μ''. We extend the solid path containing μ by concatenating a new node ν before node μ. We make Π'' the left edge-path of ν, and paths Π' and Π'' the left and right edge-paths of node μ. Finally, we call $conceal(\nu)$ to restore the size invariant.

Operation $Contract(\mu)$ is performed similarly. We begin by calling $expose(\mu)$ which results in the left edge-path of μ, $\Pi_\ell(\mu)$ containing all the edges from children to μ. Let node ν be the parent of μ. We remove node μ. We set $\Pi_\ell(\nu)$ to the concatenation of $\Pi_\ell(\nu)$, $\Pi_\ell(\mu)$, and $\Pi_r(\nu)$. We set $\Pi_r(\nu)$ to the empty path. Finally, we call $conceal(\nu)$ to restore the size invariant.

Theorem 3 *Let $\mathcal{Z} = (\mathcal{N}, \mathcal{P}, \mathcal{R}, F)$ be a tree attribute system, and $\mathcal{S}$ be a collection of path-selection and tree-selection functions. There is a fully dynamic data structure for maintaining $\mathcal{Z}$ and $\mathcal{S}$ over a set of unbounded degree trees with the following performance:*

1. *a tree of size n uses $O(n)$ space;*
2. *operations* MakeTree *and* DeleteTree *take each $O(1)$ time;*
3. *operations* Sibling *takes $O(\log n + k)$ time to return k nodes.*
4. *operations* Evaluate, LocalUpdate, GlobalUpdate, Link, Cut, Evert, Reflect, Cycle, LCA, PathFind, TreeFind, Parent, Expand, *and* Contract *take each $O(\log n)$ time, where n is the total size of the trees involved in the operation.*

3 Expression Trees

3.1 Linear Expression Trees

In this section, we consider the dynamic evaluation of linear expressions, which extends and generalizes the dynamic expression trees of [9]. We show how to solve this problem using a

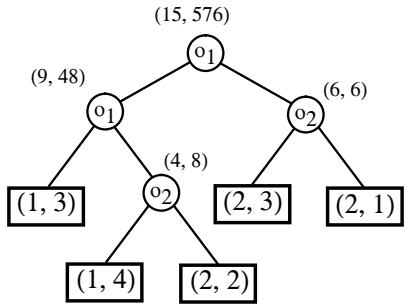

Figure 5: A linear expression tree with two binary linear operators. For $x = (x_1, x_2)$ and $y = (y_1, y_2)$ operators are defined to be: $o_1(x, y) = (x_1 + y_2, 2y_1y_2)$ and $o_2(x, y) = (2x_1 + x_1y_2, x_2y_1)$. The unary operators associated with the edges are the identity operators.

tree attribute system. Let $\mathcal{S}$ be a semiring with binary operators $\oplus$ and $\otimes$, such that $\otimes$ is distributive with respect to $\oplus$. We assume that elements of $\mathcal{S}$ can be stored in $O(1)$ space and binary operations can be performed in $O(1)$ time. The extension to the general case is straightforward. We consider operations in $\mathcal{S}^r$, the space of r-tuples of $\mathcal{S}$ for some fixed r. Let $x = (x_1, \ldots, x_r)$ and $y = (y_1, \ldots, y_r) \in \mathcal{S}^r$. A *binary linear operator* $\odot : \mathcal{S}^r \times \mathcal{S}^r \to \mathcal{S}^r$ is defined by

$$(x \odot y)_i = x^T A_i y \oplus b_i^T x \oplus c_i^T y \oplus d_i$$

for $i = 1, \cdots, r$, where $A_i \in \mathcal{S}^r \times \mathcal{S}^r$, $b_i, c_i \in \mathcal{S}^r$, and $d_i \in \mathcal{S}$. A *unary linear operator* $\nabla : \mathcal{S}^r \to \mathcal{S}^r$ is defined by

$$\nabla x = Ax \oplus b$$

where $A \in \mathcal{S}^r \times \mathcal{S}^r$, and $b \in \mathcal{S}^r$. (In the above definitions we have used standard linear algebra notation, and convention that xy denotes $x \otimes y$.) Note that, given a binary linear operator $\odot$, for any fixed $\bar{y}$, $x \odot \bar{y}$ is a unary linear operator. Also, given a unary linear operator ∇, there exists a binary linear operator $\odot$ such that $\nabla x = x \odot 0$.

A *linear expression* $\mathcal{E}$ over $\mathcal{S}^r$ is an expression with variables in $\mathcal{S}^r$ involving binary and unary linear operators. A linear expression is represented by a tree T called a linear expression tree, such that the internal nodes are the operators and the leaves are the variables of $\mathcal{E}$ (see an example in Fig. 5). We associate a binary linear operator with each internal node of T. If μ is an internal node associated with a binary linear operator, but with only one child, then μ takes the value of its child.

Example 6 Fig. 5 shows an example of linear expression and its associated binary tree.

We allow unbounded degree linear expression trees. Suppose $\odot$ is a binary linear operator. The unbounded version of $\odot$ takes a variable and unbounded number of operands. We write $\odot(x_1, \ldots, x_d)$ to indicate applying $\odot$ to values $x_1, \ldots, x_d$. If $d = 1$ then the value of $\odot(x_1)$ is $x_1 \odot 0$. Otherwise, the value of $\odot(x_1, \ldots, x_d)$ is $((x_1 \odot x_2) \ldots) \odot x_d$. An *unbounded degree linear expression* is an expression containing unbounded versions of binary linear operators.

An unbounded linear expression is represented by an unbounded linear expression tree. In order to keep the value invariant each binary linear operator must be associative.

We consider a dynamic environment where linear expressions are manipulated by changing the values of the variables and by linking and cutting the corresponding trees. We also allow a somehow unorthodox modification of an expression $\mathcal{E}$ by everting its tree T.

In order to give a meaning to the eversion of T, we require the following. Suppose μ is a node that is not the root of T. Everting at node μ results in μ having an additional child. We do not allow eversion at a leaf of a linear expression tree. In order to give meaning to eversion at a node with 1 child, each internal node calculates a binary linear operator. A unary linear operator is simulated by a binary linear operator with only one child. Finally, we allow the root ρ of T to have three children. If $\odot$ is the binary linear operator at ρ, then the value of ρ is calculated as in the unbounded version of $\odot$. As with general unbounded linear operators, we require that if μ is a node of T with 2 children, then eversion is allowed at μ only if the operator $\odot$ at μ is associative.

Given a linear expression tree T we will show that we can maintain a tree attribute system to find the values of subexpressions of T.

Lemma 3 *Suppose $\mathcal{E}$ is a linear expression. The dependence of the value y of $\mathcal{E}$ on a single variable x can be expressed by a unary linear operator $y = Ax \oplus b$.*

Proof: The proof is by induction on the level of the subexpression containing x. If y is x, then the unary linear operator is $y = x$. Now suppose $y = y' \odot y''$ and x is used in the calculation of y'. By the inductive hypothesis, we have $y' = A'x \oplus b'$ for some matrix A' and vector b'. When we fix the value of y'', we get by the earlier observation matrix A'' and vector b'' such that $y = A''y' \oplus b''$. By composition, we have $y = A''A'x \oplus (A''b' \oplus b'')$ □

We call *transfer pair* a (matrix, vector) pair (A, b) that by Lemma 3 characterizes the dependency of a linear expression on a variable. In particular, given a (solid or edge) path Π of T, we denote with $P(\Pi)$ the transfer pair associated with the dependency of the value of $tail(\Pi)$ from the value of $head(\Pi)$.

Lemma 4 *Let $\mathcal{E}$ be a linear expression represented by a tree T, and let $x(\mu)$ be the value of the subexpression given by the subtree of T rooted at μ. Then the transfer pairs $P(\Pi)$ of the solid or edge paths Π of T form a path attribute system, and there exists a tree attribute system that supports tree attribute $x(\mu)$ by means of the path attribute $P(\Pi)$.*

Proof: Suppose path Π of T is the concatenation of paths Π' and Π'' with transfer pairs (A', b') and (A'', b''), respectively. The value of $head(\Pi'')$ is obtained from the values of the children of $head(\Pi'')$ by means of a linear operator; hence $x(head(\Pi'')) = Cx(tail(\Pi')) \oplus d$. By Lemma 3 we have $x(head(\Pi)) = A''(C(A'x(tail(\Pi)) \oplus b') \oplus d) \oplus b''$. Thus, $x(head(\Pi)) = A''CA'x(tail(\Pi)) \oplus A''(Cb' \oplus d) \oplus b''$, so that the transfer pair $P(\Pi) = (A, b)$ is given by $A = A''CA'$ and $b = A''(Cb' \oplus d) \oplus b''$. This shows that the transfer pairs form a path attribute system. Clearly, transfer pair $P(\Pi)$ allows to compute in $O(1)$ time $x(tail(\Pi))$ from $x(head(\Pi))$. Also, the transfer pairs depend on an extended node attribute set that for a node μ of a path Π consists of the operator of μ and the value $x(\mu')$ of the child μ' of μ not in Π. Hence, the triplet (O, P, x), where $O(\mu)$ denotes the operator of node μ, is a tree attribute system, and the lemma follows from Theorem 2. □

Theorem 4 *Let $\mathcal{S}$ be a semiring and r an integer. There is a fully dynamic data structure for maintaining a set of expressions with (unary and binary) linear operators over $\mathcal{S}^r$ with the following performance:*

1. *an expression of size n uses $O(n)$ space;*
2. *operations* MakeTree *and* DeleteTree *take each $O(1)$ time;*
3. *operations* Evaluate, LocalUpdate, Link, Cut, Evert, Reflect, Cycle, PathFind, TreeFind, Parent, *and* Sibling, *take each $O(\log n)$ time, where n is the total size of the trees involved in the operation.*

If we restrict ourselves to associative binary linear operators, then by lemma 2, lemma 4, and Theorem 2, we get:

Theorem 5 *Let $\mathcal{S}$ be a semiring and r an integer. There is a fully dynamic data structure for maintaining a set of unbounded linear expressions with associative binary linear operators over $\mathcal{S}^r$ with the following performance:*

1. *an expression of size n uses $O(n)$ space;*
2. *operations* MakeTree *and* DeleteTree *take each $O(1)$ time;*
3. *operations* Sibling *takes $O(\log n + k)$ time to return k nodes.*
4. *operations* Evaluate, LocalUpdate, Link, Cut, Evert, Reflect, Cycle, LCA, PathFind, TreeFind, Parent, Expand, *and* Contract *take each $O(\log n)$ time, where n is the total size of the trees involved in the operation.*

3.2 Linear Attribute Grammars

In this section, we consider the dynamic evaluation of linear attribute grammars and show that this problem can also be solved using a tree attribute system.

An *attribute grammar* T is a rooted, ordered tree such that each node μ of T stores an r-tuple of attributes $x(\mu) = (x_1(\mu), \ldots, x_r(\mu))$. *Synthesized* attributes are calculated from the children of μ. *Inherited* attributes are calculated from the attributes of the parent of μ, and the synthesized attributes of μ and its siblings. We assume, without loss of generality, that the synthesized attributes of μ precede the inherited attributes of μ. A *binary attribute grammar* is an attribute grammar on a binary tree. We consider attribute grammars where the values of each attribute are taken from a semiring $\mathcal{S}$.

Attribute grammars [20] are much studied and have applications in areas such as language based editors [3,26], and VLSI design [18]. Incremental algorithms for attribute grammars have been presented [17,25] which give a $O(a \cdot t)$ running time per operation where a is the number of attributes affected by a change and t is the time to make a single change. In general, $a = O(n)$ for an attribute grammar of size n. Incremental evaluation of general circuits is studied in [1].

The precedence graph for an attribute grammar T is a digraph consisting of a vertex for each attribute $x_i(\mu)$ and a directed edge from $x_i(\mu)$ to $x_j(\nu)$ if the value $x_i(\mu)$ is used in the calculation of $x_j(\nu)$.

For a node μ of attribute grammar T, let $s(\mu) = (s_1(\mu), \ldots s_r(\mu))$ and $h(\mu) = (h_1(\mu), \ldots h_r(\mu))$ be the r-tuples defined by:

$$s(\mu) = \begin{cases} x_j(\mu) & \text{if } x_j(\mu) \text{ is synthesized} \\ 0 & \text{if } x_j(\mu) \text{ is inherited} \end{cases}$$

$$h(\mu) = \begin{cases} x_j(\mu) & \text{if } x_j(\mu) \text{ is inherited} \\ 0 & \text{if } x_j(\mu) \text{ is synthesized} \end{cases}$$

where $1 \leq j \leq r$.

Suppose node μ has $p < r$ synthesized attributes. Then $s^R(\mu) = (x_1(\mu), \ldots, x_p(\mu))$, and $h^R(\mu) = (x_{p+1}(\mu), \ldots, x_r(\mu))$.

Definition 18 *Consider a node μ of a binary attribute grammar T. Let node ν be the parent of μ, node μ' be the sibling of μ, and nodes μ_1 and μ_2 be the children of μ. A* bidirectional binary operator $\odot = (\odot_S, \odot_I, \odot_\ell, \odot_r)$ *is a 4-tuple of binary linear operators such that:*

$$s(\mu) = x(\mu_1) \odot_S x(\mu_2)$$

and, if μ is the left child of ν:

$$h(\mu) = s(\mu) \odot_I (x(\nu) \odot_r s(\mu'))$$

otherwise, if μ is the right child of ν:

$$h(\mu) = s(\mu) \odot_I (x(\nu) \odot_\ell s(\mu'))$$

Definition 19 *A* binary linear attribute grammar *T is a binary attribute grammar such that:*

- *The value of all attributes come from a semiring $\mathcal{S}$.*
- *Each node μ of T has an associated bidirectional operator to calculate the value $x(\mu)$.*
- *The precedence graph of G is acyclic.*
- *The dependence of the value an attribute $x_i(\mu)$ on the value of any other attribute $x_j(\nu)$ is linear. That is, we can express the dependence in the form $x_i(\mu) = ax_j(\nu) + b$ where a and b are members of $\mathcal{S}$.*

A (binary) linear expression tree is a (binary) linear attribute grammar without inherited attributes.

Suppose G is the precedent graph of a linear attribute grammar T and $x_i(\nu_1)$ and $x_j(\nu_2)$ are two nodes of G. Suppose π' and π'' are two paths in G from $x_i(\nu_1)$ to $x_j(\nu_2)$ that only meet at their endpoints. Let $x_h(\mu')$ and $x_k(\mu'')$ be the predecessors of $x_j(\nu_2)$ on π' and π''. Then the attributes $x_h(\mu')$ and $x_k(\mu'')$ are not multiplied by the operation performed at ν_2.

Suppose $x = (x_1, \ldots, x_h)$ is a vector. The power vector $power(x)$ is a vector with 2^h entries $(1, x_1, \ldots, x_h, x_1x_2, \ldots, x_{h-1}x_h, \ldots, x_1 \cdots x_h)$.

Lemma 5 *Suppose G is a precedence graph such that all dependencies are linear. Consider attributes $y_0, \ldots, y_f$ of G. The dependence of the value of y_0 on the values of $y_1, \ldots, y_f$ can be expressed as $y_0 = a \cdot \text{power}((y_1, \ldots, y_f))$ for some vector a with 2^f components.*

Proof: Let G be the precedence graph of T. Our proof is by induction on the length of a longest directed path from some y_i $(1 \leq i \leq f)$ to y_0 in G. If the length is 0, then all y_i $(0 \leq i \leq f)$ are identical, and the lemma follows with $a = (0, 1, 0, \ldots)$.

Otherwise, let $(\nu_1, \ldots, \nu_k)$ be the attributes of T such that for each ν_j $(1 \leq j \leq k)$ there is an edge from ν_j to y_0 and a path of length greater than or equal to 0 from some y_i to ν_j. Then for each ν_j there is a vector a_j such that $\nu_j = a_j \cdot power((y_1, \ldots, y_j))$. Additionally, there is a vector b such that $y_0 = b \cdot power((\nu_1, \ldots, \nu_k))$. By substitution, we get a polynomial formula for y_0 in terms of $y_1, \ldots, y_j$. No term of this formula can contain y_j^g, $g \geq 2$, for $1 \leq j \leq k$ since this would violate the linear dependency of attributes in a linear attribute grammar. □

Corollary 1 *Suppose μ is a node of linear attribute grammar T. The dependence of a synthesized attribute $x_i(\mu)$ on the inherited attributes $h^R(\mu)$ can be expressed as $x_i(\mu) = a \cdot \text{power}(h^R(\mu))$.*

The vector a depends only on the calculations of attribute values performed at nodes in the subtree rooted at μ, excluding the calculation of the inherited attributes of μ.

Corollary 2 *Suppose Π is a path in linear attribute grammar T with head σ and tail τ. Let y be the vector consisting of the attributes $s^R(\sigma)$ followed by the attributes $h^R(\tau)$. The dependence of an inherited attribute $x_i(\sigma)$ on the attributes of y can be expressed as $x_i(\sigma) = a \cdot \text{power}(y)$ for some vector a. Similarly the dependence of synthesized attribute $x_j(\tau)$ on the attributes of x can be expressed as $x_j(\tau) = b \cdot \text{power}(y)$ for some vector b.*

The vectors a and b depends on the calculations of attribute values performed at nodes in the subtree rooted at τ excluding the subtree rooted at σ.

Suppose we are given a directed acyclic graph G. Graph G is an *evaluation graph* if each vertex v of G calculates a value $value(v)$ by an equation $value(v) = a \cdot power(y)$, where y is the vector of values of vertices u with edges in G from u to v. A *topological sort* of G orders the vertices of G such that if there is an edge from vertex u to vertex v in G, then u precedes v in the ordering. A *simple evaluation scheme* calculates the values for each vertex of G in topological order. In this way, the input values are available when calculating the value of a vertex.

For a node μ in linear attribute grammar T, corollary 1 shows that we can summarize the dependence of the attributes of $s^R(\mu)$ on the attributes of $h^R(\mu)$ as an evaluation graph $summarygraph(\mu)$. There is a vertex of $summarygraph(\mu)$ for each attribute of μ. Edges are directed from attributes of $s^R(\mu)$ to attributes of $h^R(\mu)$. Similarly, for a path Π in T with head σ and tail τ, corollary 2 shows that we can summarize the dependence of the attributes of $h^R(\sigma)$ and $s^R(\tau)$ on the attributes of $s^R(\sigma)$ and $h^R(\tau)$ as an evaluation graph $summarygraph(\Pi)$. There is a vertex of $summarygraph(\Pi)$ for each attribute of σ and τ. Edges are directed from attributes of $s^R(\sigma)$ and $h^R(\tau)$ to attributes of $h^R(\sigma)$ and $s^R(\tau)$. Notice that for any node μ and any path Π, evaluation graphs $summarygraph(\mu)$ and $summarygraph(\Pi)$ will be of constant size.

Lemma 6 *Let T be a linear attribute grammar, $R(\mu)$ be the graph* summarygraph(μ)*, $N(\mu)$ be the bidirectional binary operator $\odot$, and $P(\Pi)$ be the directed graph* summarygraph(Π)*. Then there exists a tree attribute system that supports tree attribute $R(\mu)$ by means of the path attribute $P(\Pi)$.*

Proof: Suppose Π is a path in T from node σ to node τ. Evaluation graph $summarygraph(\Pi)$ gives the dependencies of the attributes of $h^R(\sigma)$ and $s^R(\tau)$ on the attributes of $s^R(\sigma)$ and $h^R(\tau)$. Evaluation graph $summarygraph(\sigma)$ gives the dependencies of the attributes of $s^R(\sigma)$

on the attributes of $h^R(\sigma)$. We combine these two graphs and perform the simple evaluation scheme to get the dependencies of the attributes of $s^R(\tau)$ on the attributes of $h^R(\tau)$, which is $summarygraph(\tau)$. All this can be done in constant time, since these graphs are of bounded size.

To show that $P(\Pi)$ can be maintained in a path attribute system, consider path Π' from head σ' to τ' and path Π'' from head σ'' to τ'' with node σ'' the parent of node τ'. Suppose node μ is the sibling of node τ'. The evaluation graph $summarygraph(\mu)$ will be part of the extended node attribute set $\tilde{\mathcal{N}}(\sigma'')$.

Suppose path Π is the concatenation of paths Π' and Π''. To find $summarygraph(\Pi)$, we form an evaluation graph H with vertices for each of the attributes of $x(\sigma')$, $x(\tau')$, $x(\sigma'')$, $x(\tau'')$, and $x(\mu)$. The edges of H are the edges of $summarygraph(\Pi')$, $summarygraph(\Pi'')$, $summarygraph(\mu)$, the dependencies for calculating $s^R(\sigma'')$ from $x(\tau')$ and $x(\mu)$, and the dependencies for calculating $h^R(\tau')$ from $s^R(\tau')$, $s^R(\mu)$, and $x(\sigma'')$. We perform the simple evaluation scheme on H to determine the dependencies of the attributes of $h^R(\sigma')$ and $s^R(\tau'')$ on the attributes of $s^R(\sigma')$ and $h^R(\tau'')$, which is $summarygraph(\Pi)$ (see Fig. 6). All this can be done in constant time, since these graphs are of bounded size. □

For a node μ of linear attribute grammar T, the value of $Evaluate(\mu)$ is the evaluation graph $summarygraph(\mu)$. In order to return the actual values of $x(\mu)$, we present the following operation for linear attribute grammars.

- *EvaluateAttributes*(*node* μ) — Returns the values of the attributes $x(\mu)$.

We implement operation $EvaluateAttributes(\mu)$ as follows. First we call $Evaluate(\mu)$ to get evaluation graph $summarygraph(\mu)$. We then call $expose(\mu)$ to get path Π from μ to root ρ. By combining the evaluation graphs $summarygraph(\Pi)$ and $summarygraph(\mu)$, and performing the simple evaluation scheme, we get the dependencies of $x(\mu)$ on $h^R(\rho)$. We then calculate the actual values of $x(\mu)$ by plugging in the default inherited value for the root ρ of T. We finish by calling $conceal(\mu)$ to restore the size invariant.

Theorem 6 *Let $\mathcal{S}$ be a semiring and r an integer. There is a fully dynamic data structure for maintaining a set of linear attribute grammars over $\mathcal{S}^r$ with the following performance:*

1. *an linear attribute grammar of size n uses $O(n)$ space;*
2. *operations* MakeTree *and* DeleteTree *take each $O(1)$ time;*
3. *operations* EvaluateAttributes, LocalUpdate, Link, Cut, Evert, Reflect, Cycle, *LCA*, PathFind, TreeFind, Parent, *and* Sibling, *take each $O(\log n)$ time, where n is the total size of the linear attribute grammars involved in the operation.*

A binary linear operator is an *unbounded linear operator* if it keeps the expansion invariant. That is, if μ is a node of linear attribute grammar T, then $summarygraph(\mu)$ is unchanged after expanding any children of μ. Note that the actual value of attributes may change as as a result of an *Expand* operation.

Consider the following example. Given a node μ in linear attribute grammar T, let $depth(\mu)$ be the depth of node μ in T. The attribute $depth(\mu)$ is inherited. The value $depth(\mu)$ is either 0, if μ is the root of T, or $depth(\nu)+1$, where node ν is the parent of μ. If μ is an internal node and μ' is the leftmost child of μ. Let $synthesizeddepth(\mu) = depth(\mu')-1$. Therefore, we get the relation $synthesizeddepth(\mu) = depth(\mu)$, which remains true after any

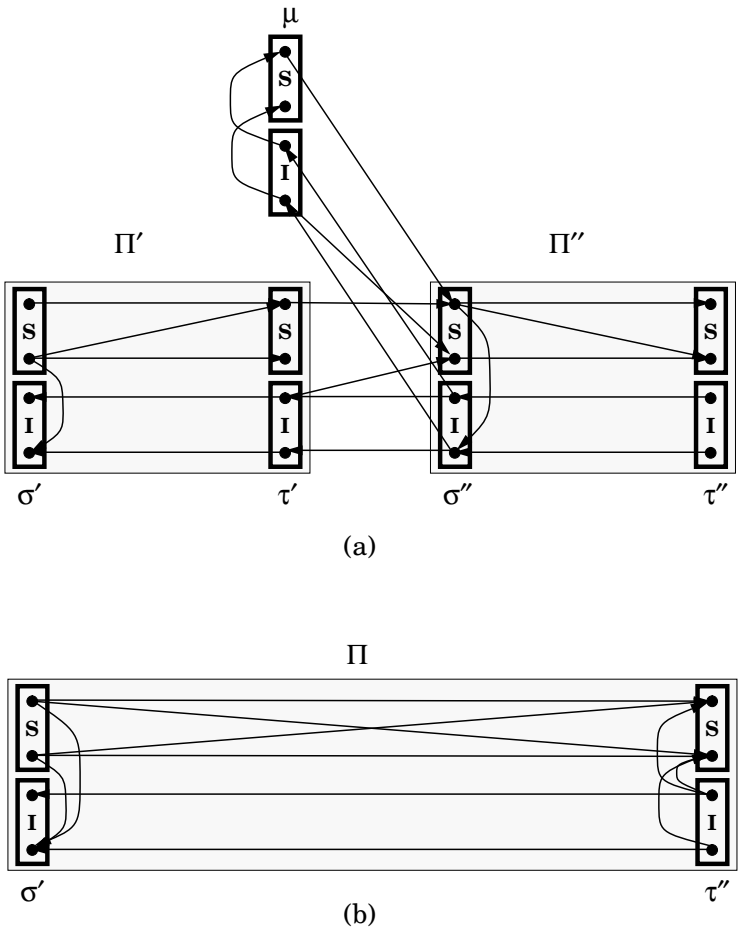

Figure 6: An example of the concatenation of summary graphs as described in lemma 6. (a) The combined evaluation graph. (b) The resulting summary graph

expansion of a node of T. Yet, when expanding children of μ, we increment the depth of the descendants of the expanded nodes.

An *unbounded linear attribute grammar* is a linear attribute grammar with unbounded degree nodes such that each node calculates an unbounded linear operator.

By by lemma 2, lemma 6, and Theorem 2, we get:

Theorem 7 *Let $\mathcal{S}$ be a semiring and r an integer. There is a fully dynamic data structure for maintaining a set of unbounded linear attribute grammars over $\mathcal{S}^r$ with the following performance:*

1. *an linear attribute grammar of size n uses $O(n)$ space;*
2. *operations* MakeTree *and* DeleteTree *take each $O(1)$ time;*
3. *operations* Sibling *takes $O(\log n + k)$ time to return k nodes.*
4. *operations* EvaluateAttributes, LocalUpdate, Link, Cut, Evert, Reflect, Cycle, LCA PathFind, TreeFind, Parent, Expand, *and* Contract *take each $O(\log n)$ time, where n is the total size of the linear attribute grammars involved in the operation.*

4 Applications

In this section we show applications of tree attribute systems, linear expression trees, and linear attribute grammars. We show how a number of existing data structures can be expressed as tree attribute systems. We also present new applications using our techniques.

4.1 Tree Structures for Graph Problems

A number of existing data structures can be expressed as tree attribute systems. This list includes dynamic trees, dynamic expression trees, and edge ordered trees. In this section, we review these data structures, and existing applications for tree attribute systems, linear expression trees, and linear attribute grammars.

Dynamic trees [28] are unordered, unbounded, rooted trees supporting operations *Link*, *Cut* and *Evert*. Each node μ stores a globally updatable cost associated with the edge from μ to its parent. Global updates are only performed on the path from a node to the root. A specialized *PathFind* operator is included to find the minimum cost edge on a path. Dynamic trees were first presented as an internal data structure in several (sequential) maximum flow algorithms.

Dynamic expression trees [9] are a restricted form of linear expression trees. Values are scalars taken from a semi-ring $\mathcal{S}$. Operations are the $+$ and $\times$ operators of semi-ring $\mathcal{S}$. The major presented algorithm for dynamic expression trees is fully dynamic algorithms for maximum flow and shortest path in series-parallel digraphs.

In [8], linear expression trees are used to present fully dynamic algorithms for a large number of problems on tree-width two graphs, the class of graphs where each biconnected component is a biconnected series-parallel graph. Using linear expression trees augmented with additional node, path, and tree attributes, fully dynamic algorithms are given for queries such as the size of a vertex cover, dominating set, chromatic number, maximum independent set, connectible pairs shortest path, and connectible pairs maximum flow. Queries are performed in $O(\log m)$ time, while updates are performed in $O(\log^2 m)$ time for a graph with m edges.

A number of data structures have been presented which maintain the recursive decomposition of graphs into subgraphs with constant size overlap. SPQ-trees [9] and SPQC-trees are used as a fully dynamic data structure to maintain a decomposition of series-parallel digraphs and tree-width two graphs. SPQR-trees [11] are used to maintain the tri-connected components of an initially biconnected graph. FWRT-trees [19] used to maintain the 4-connected components of an initially tri-connected graph. These last two data structures have a specialized *PathFind* operation to find the closest ancestor of a given type.

Edge-ordered trees [12] are ordered, unbounded, rooted trees supporting operations *Link*, *Cut*, *Expand*, *Contract*, *Cycle* and *Evert*. Each node μ stores a locally updatable cost associated with the edge from μ to its parent. A specialized *PathFind* operator is included to find the minimum cost edge on a path. Edge-ordered trees were presented as the data structure for a fully dynamic algorithm to maintain the minimum or maximum spanning tree of planar graphs.

In [7], simplified linear attribute grammars are used to maintain drawings of series-parallel graphs and trees. Queries include location of vertices and edges of the drawing,

point location, and window queries. Queries take $O(k \log m)$, where k is the number of vertices and edges returned, and m is the number of edges of the tree or series-parallel graph.

4.2 Summation Heaps

In this section we present a generalization of heaps called summation heaps. Recall that in the tree implementation of a heap, the structure of the tree is arbitrary (see example 5). Heap elements are stored at the leaves of the tree. Internal nodes are only used to direct us when selecting the minimum value of a heap.

In summation heaps, the structure of the heap is important. We extend costs to include internal nodes. A summation heap is an expression tree over operators + and min. The cost of a node μ is the value of the subexpression rooted at μ.

We replace operation *FindMin* for heaps, with the following operation.

- *FindMinNode*(*node* ν) — return a minimum cost node μ in the subtree rooted at ν.

To implement operation *FindMinNode*, we calculate two values for each node μ:

- $cost(\mu)$ — the cost of node μ.
- $mincost(\mu)$ — the minimum cost of a node in the subtree rooted at μ.

These values are calculated as linear expressions The value of $cost(\mu)$ is defined to be the result of a linear expression. The value of $mincost(\mu)$ is the minimum of $cost(\mu)$ and *mincost* for the children of μ. Operation $FindMin(\mu)$ is performed by calling operation $TreeFind(\mu, S_4)$, where S_4 is the selection function of example 5.

Theorem 8 *There exists a fully dynamic $O(n)$ space data structure for maintaining a collection of summation heaps of total size n, such that a query or update operation takes time $O(\log n)$*

4.3 Blocking Heaps

In this section we present another generalization of heaps called blocking heaps, which are used in the implementation of the dynamic algorithm for tree-width two graphs.

Suppose $\mathcal{S}$ is a totally-ordered set. We consider operations in $\mathcal{S}^r$, the space of r-tuples of $\mathcal{S}$ for some fixed r. Let $x = (x_1, \ldots, x_r)$ and $y = (y_1, \ldots, y_r) \in \mathcal{S}^r$. A *binary minimization operator* $\odot : \mathcal{S}^r \times \mathcal{S}^r \to \mathcal{S}^r$ is defined such that each component $(x \odot y)_i$ calculates the minimum of some subset of the components of x and some subset of the components of y. Given a component μ_i at node μ, a *source* of μ_i is a component λ_j at leaf λ such that reducing the value of λ_j, reduces the value of μ_i.

Consider the following operation on a tree associated with an expression of binary minimization operators:

- *FindSources*(*node* μ) — return a source for each of the components of μ.

Definition 20 *A* blocking heap *is a linear expression tree of binary minimization operators which also supports operation* FindSources.

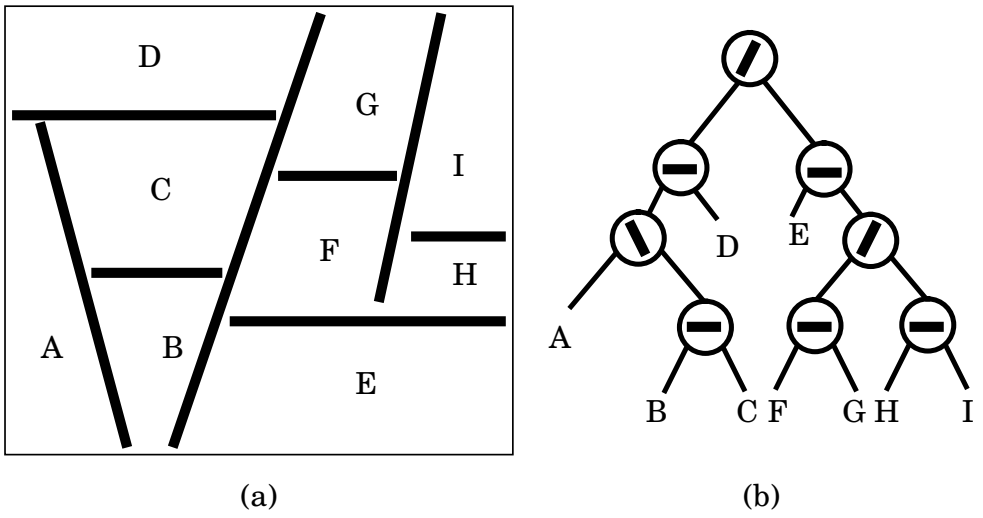

Figure 7: (a) A binary space partition of the plane into trapezoids. (b) The associated decomposition tree.

Suppose T is a blocking heap. For each path Π in T from node σ to node τ, we keep the following boolean table:

- *depends*(Π, i, j) — **true** if there is a path in the dependency graph of T from σ_i to τ_j.

Operation *FindSources*(μ) is performed by calling operation *TreeFind*(μ, S_5, i), for each $1 \leq i \leq r$, where S_5 is the following tree selection function:

Selection Function S_5 *Suppose* Π *is a path, node* ζ *the root of the path tree* $\mathcal{B}_\Pi$*, and nodes* ζ' *and* ζ'' *the left and right children of* ζ*. Let* τ' *be* tail(ζ')*. Use the transfer pairs at* ζ *and* ζ'*, the binary minimization operator at* head(ζ'')*, and the* depends *table for* ζ'' *to determine if there is a component* τ'_j *at* τ' *such that the value of* τ'_j *is the same as* μ_i*, and* μ_i *depends on* τ_j*. If so, return* ζ'*, else return* ζ''*.*

Theorem 9 *There exists a fully dynamic* $O(n)$ *space data structure for maintaining a collection of blocking heaps of total size* n*, such that a query or update operation takes time* $O(\log n)$

4.4 Point Location in Unbalanced Binary Space Partition Trees

A binary space partition is a binary tree where each node corresponds to a connected region of the space (or plane) and to a geometric object that bipartitions such a region (see Fig. 7). E.g. a node corresponds to a trapezoid and to a segment that divides it into two trapezoids (for details see [22]). Each leaf λ of binary space partition T represents a region referred to as *region*(λ). If the discrimination of a point with respect to the dividing object (i.e., figuring out which of the two children regions contains the point) can be done in polylog time (usually O(1) time), then one can do point location in polylog time. Previous approaches

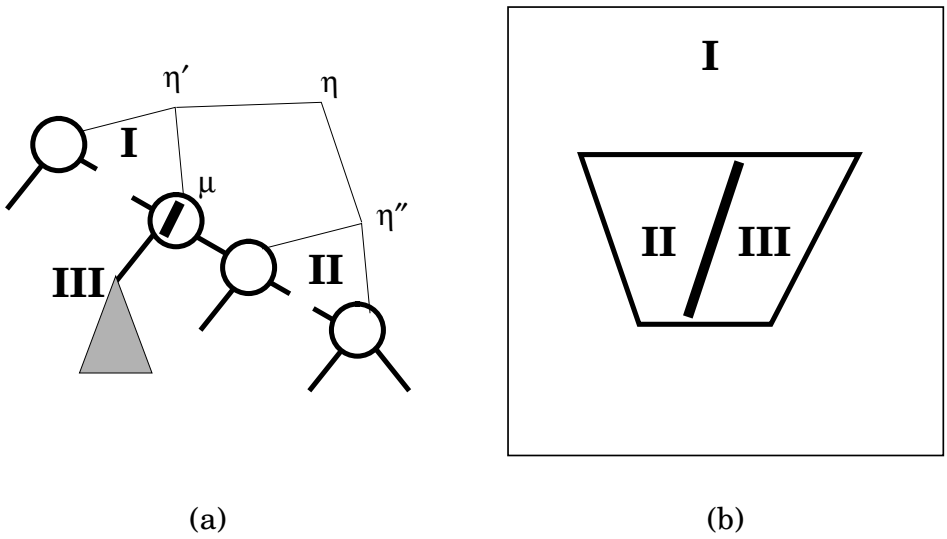

Figure 8: (a) A solid path in a decomposition tree; and (b) some regions of a binary space partition. For a point p, consider selection node $\zeta = S_6(\eta, p)$. If $\zeta = \eta''$ then p is in regions I or III. If $\zeta = \eta'$ then p is in regions II. If node $\mu = \mathit{find}(\mathrm{II}, S_6, p)$, then p is in region III.

try to keep the binary space partition balanced [24,27]. With a selection function one can use unbalanced partitions. A specialized version of this technique was introduced in [4].

We support operations on a collection of binary space partitions:

- *Region*(*node* ρ; *point* p) — Return the region containing point p in the binary space partition rooted at node ρ.
- *MakeBSP*() — Create a new elementary binary space partition with no partitions.
- *DeleteBSP*(*node* λ) — Remove the elementary binary space partition represented by node λ.
- *Divide*(*node* λ, *partition* P) — Use partition P to divide the region represented by leaf λ.
- *Compose*(*node* λ, ρ) — Partition the region represented by leaf λ using the binary space partition with root ρ.
- *Decompose*(*node* μ) — Replace the sub-partition represented by node μ with a single region.

Operations *MakeBSP* and *delbsp* correspond to operations *MakeTree* and *DeleteTree* respectively. Operations $\mathit{Compose}(\lambda, \rho)$ and $\mathit{Decompose}(\mu)$ each are implemented with a constant number of *Link* and *Cut* operations.

Operation $\mathit{Divide}(\lambda, P)$ consists of replacing node λ with a node μ representing partition P. Node μ has two children representing the regions created by partition P.

We implement operation $\mathit{Region}(\rho, p)$ using the following selection function, which takes a point p as an argument.

Selection Function S_6 *Suppose η is the root of a path tree with children η' and η''. If p is*

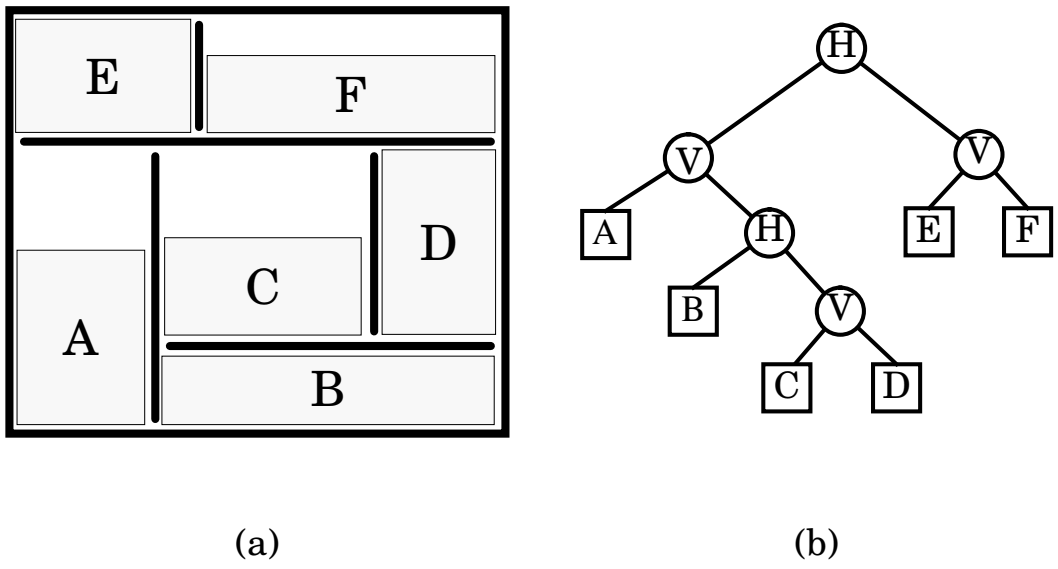

Figure 9: (a) A slicing floorplan; and (b) its associated parse tree.

contained in region(tail(η')), *then return* η', *else return* η''.

Operation *TreeFind*(ρ, S_6, p) returns a leaf λ of T such that point p is contained in *region*(λ) (see Fig. 8). We then return either *region*(λ), or if p is on the boundary of *region*(λ), the portion of the boundary containing p.

Theorem 10 *There is a fully dynamic data structure for maintaining a collection of unbalanced binary space partitions with the following performance:*

1. *operations* MakeBSP *and* DeleteBSP *take each* $O(q_1(n))$ *time, where* $q_1(n)$ *is the time required to build a single partition;*
2. *operations* Compose, Decompose, *and* Divide *take each* $O(q_1(n) \cdot \log n)$ *time;*
3. *operation* Region *takes* $O(q_2(n) \cdot \log n)$ *time, where* $q_2(n)$ *is the time to perform point location in a single region. The value* n *is the total size of the binary space partitions involved in each operation.*

The trapezoid method for planar point location described in [6] is an example of an existing method that can be expressed as a selection function.

4.5 Slicing Floorplan Compaction

Linear attribute grammars have immediate application to the problem of compacting *slicing floorplans*, a layout technique widely used in VLSI (see, e.g., [30]). A slicing floorplan is either a rectangle (called basic rectangle), or is the union of two slicing floorplans that share a horizontal side (called horizontal slice) or a vertical side (called vertical slice). Each basic rectangle r has a minimum width w_r and a minimum height h_r, which describe the dimensions of a circuit module to be placed inside r.

An important problem in VLSI layout is determining the location of basic rectangles

Value	*Vertical*	*Horizontal*
$height(\mu)$	$\max(height(\mu_1), height(\mu_2))$	$height(\mu_1) + height(\mu_2)$
$width(\mu)$	$width(\mu_1) + width(\mu_2)$	$\max(width(\mu_1), width(\mu_2))$
$xpos(\mu_1)$	$xpos(\mu)$	$xpos(\mu)$
$ypos(\mu_1)$	$ypos(\mu)$	$ypos(\mu)$
$xpos(\mu_2)$	$xpos(\mu) + width(\mu_1)$	$xpos(\mu)$
$ypos(\mu_2)$	$ypos(\mu)$	$ypos(\mu) + height(\mu_1)$

Table 3 The equations to calculate the compaction of floorplans for a node μ and its children μ_1 and μ_2. The choice of equations depend on if μ represents a vertical or horizontal slice.

while minimizing the area of the slicing floorplan, subject to the above constraints on the height and width of the basic rectangles. (For full details, see [30].) This problem can be solved using a linear attribute grammar T. Leaves of T represent basic rectangles, and internal nodes represent horizontal or vertical slices (see Fig. 9). The *reference point* of a sub-floorplan is its bottom-left corner.

We keep the following synthesized attributes at each node μ of linear attribute grammar T with root ρ:

- $height(\mu)$ — The height of the sub-floorplan represented by node μ.
- $width(\mu)$ — The width of the sub-floorplan represented by node μ.

We also keep the following inherited attributes:

- $xpos(\mu)$ — The x-coordinate of the reference point of the sub-floorplan represented by node μ. The value $xpos(\rho)$ is assumed to be 0.
- $ypos(\mu)$ — The y-coordinate of the reference point of the sub-floorplan represented by node μ. The value $ypos(\rho)$ is assumed to be 0.

The equations to calculate these values is shown in table 3. Using linear attribute grammars, we can solve a dynamic version of the compaction problem for slicing floorplans, where query operations ask for the minimum area of a floorplan (or sub-floorplan) or the location of a basic rectangle and update operations change the dimensions of a basic rectangle or join (or separate) two floorplans.

Theorem 11 *There exists a fully dynamic $O(n)$-space data structure for maintaining a collection of slicing floorplans of total size n, such that a query or update operation takes $O(\log n)$ worst-case time.*

References

[1] B. Alpern, R. Hoover, B. Rosen, P. Sweeney, and F.K. Zadeck, "Incremental Evaluation of Computational Circuits," *Proc. ACM-SIAM Symp. on Discrete Algorithms* (1990), 32–42.

[2] S.W. Bent, D.D. Sleator, and R.E. Tarjan, "Biased Search Trees," *SIAM J. Computing* 14 (1985), 545–568.

[3] G. M. Beshers and R. H. Campbell, "Maintained and Constructor Attributes," *Proc. ACM Symp. on Language Issues in Programming Environments, ACM SIGPLAN Notices* (1985), 34–42.

[4] Y.-J. Chiang, F.P. Preparata, and R. Tamassia, "A Unified Approach to Dynamic Point Location, Ray Shooting and Shortest Paths in Planar Maps," Dept. Computer Science, Brown Univ., Technical Report CS-92-07, 1992.

[5] Y.-J. Chiang and R. Tamassia, "Dynamic Algorithms in Computational Geometry," Dept. of Computer Science, Brown Univ., Technical Report CS-91-24, 1991.

[6] Y.-J. Chiang and R. Tamassia, "Dynamization of the Trapezoid Method for Planar Point Location," *Int. J. of Computational Geometry Applications* (to appear).

[7] R. F. Cohen, G. Di Battista, R. Tamassia, I.G. Tollis, and P. Bertolazzi, "A Framework for Dynamic Graph Drawing," *Proc. ACM Symp. on Computational Geometry* (1992).

[8] R.F. Cohen, S. Sairam, R. Tamassia, and J.S. Vitter, "Dynamic Algorithms for Bounded Tree-Width Graphs," Brown University, Manuscript, 1992.

[9] R.F. Cohen and R. Tamassia, "Dynamic Expression Trees and their Applications," *Proc. ACM-SIAM Symp. on Discrete Algorithms* (1991), 52–61.

[10] G. Di Battista and R. Tamassia, "Incremental Planarity Testing," *Proc. 30th IEEE Symp. on Foundations of Computer Science* (1989), 436–441.

[11] G. Di Battista and R. Tamassia, "On-Line Graph Algorithms with SPQR-Trees," *Automata, Languages and Programming (Proc. 17th ICALP), Lecture Notes in Computer Science* 442 (1990), 598–611.

[12] D. Eppstein, G.F. Italiano, R. Tamassia, R.E. Tarjan, J. Westbrook, and M. Yung, "Maintenance of a Minimum Spanning Forest in a Dynamic Plane Graph," *J. of Algorithms* 13 (1992), 33–54.

[13] G.N. Frederickson, "Data Structures for On-Line Updating of Minimum Spanning Trees, with Applications," *SIAM J. Computing* 14 (1985), 781–798.

[14] G.N. Frederickson, "Ambivalent Data Structures for Dynamic 2-Edge-Connectivity and k Smallest Spanning Trees," *Proc. 32th IEEE Symp. on Foundations of Computer Science* (1991).

[15] Z. Galil and G.F. Italiano, "Fully Dynamic Algorithms for Edge-Connectivity Problems," *Proc. 23th ACM Symp. on Theory of Computing* (1991), 317–327.

[16] Z. Galil and G.F. Italiano, "Maintaining Biconnected Components of Dynamic Planar Graphs," *Automata, Languages and Programming (Proc. 18th ICALP), Lecture Notes in Computer Science* (1991).

[17] L. G. Jones, "Incremental Compaction of Flat Symbolic IC Layouts," Department of Computer Science, University of Illinois, Urbana, Illinois, Technical Report No. UIUCDCS-R-87-1386, 1987.

[18] L. G. Jones and J. Simon, "Hierarchical VLSI Design Systems Based on Attribute Grammars," *Proc. 13th ACM Symp. on Principles of Programming Languages* (1986), 58–69.

[19] A. Kanevsky, R. Tamassia, J. Chen, and G. Di Battista, "On-Line Maintenance of the Four-Connected Components of a Graph," *Proc. 32th IEEE Symp. on Foundations of Computer Science* (1991), 793–801.

[20] D. E. Knuth, "Semantics of Context-Free Languages," *Mathematical Systems Theory* 2 (1968), 127–145.

[21] M. Overmars, "The Design of Dynamic Data Structures," *Lecture Notes in Computer Science* 156 (1983).

[22] M.S. Patterson and F.F. Yao, "Optimal Binary Space Partitions for Orthoganal Objects," *Proc. 1st ACM-SIAM Symp. on Discrete Algorithms* (1990), 100–106.

[23] J.A. La Poutre, "Dynamic Graph Algorithms and Data Structures," Dept. of Computer Science, University of Utrechet, Utrechet, Ph.D. Thesis, 1991.

[24] F.P. Preparata, "A New Approach to Planar Point Location," *SIAM J. Computing* 10 (1981), 473–483.

[25] T.W. Reps, *Generating Language-Based Environments*, The MIT Press, 1984.

[26] T.W. Reps and T. Teitelbaum, *The Synthesizer Generator*, Springer-Verlag, 1989.

[27] C. Schwarz, M. Smid, and J. Snoeyink, "An Optimal Algorithm for the On-line Closest-Pair Problem," *Proc. ACM Symp. on Computational Geometry* (1992).

[28] D.D. Sleator and R.E. Tarjan, "A Data Structure for Dynamic Trees," *J. Computer Systems Sciences* 24 (1983), 362–381.

[29] D.D. Sleator and R.E. Tarjan, "Self-Adjusting Binary Search Trees," *J. ACM* 32 (1985), 652–686.

[30] L. Stockmeyer, "Optimal Orientation of Cells in Slicing Floorplan Design," *Information and Control* 57 (1983), 91–101.

[31] J. Westbrook and R.E. Tarjan, "Maintaining Bridge-Connected and Biconnected Components On-Line," *Algorithmica* (1989 (to appear in Algorithmica)).