

Fully Persistent Lists with Catenation

JAMES R. DRISCOLL

Dartmouth College, Hanover, New Hampshire

DANIEL D. K. SLEATOR

Carnegie-Mellon University, Pittsburgh, Pennsylvania

AND

ROBERT E. TARJAN

Princeton University, Princeton, New Jersey, and NEC Research Institute, Princeton, New Jersey

Abstract. This paper considers the problem of representing stacks with catenation so that any stack, old or new, is available for access or update operations. This problem arises in the implementation of list-based and functional programming languages. A solution is proposed requiring constant time and space for each stack operation except catenation, which requires $O(\log \log k)$ time and space. Here k is the number of stack operations done before the catenation. All the resource bounds are amortized over the sequence of operations.

Categories and Subject Descriptors: D.1.1 [Programming Techniques]: Applicative (Functional) Programming; E.1 [Data Structures]: Lists; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*computations on discrete structures*

General Terms: Algorithms, Design, Languages, Theory

Additional Key Words and Phrases: Amortization, catenation, concatenation, data structures, functional programming, LISP, list, queue, stack

A preliminary version of this paper appeared in *Proceedings of the 2nd Annual ACM-SIAM Symposium on Discrete Algorithms* (San Francisco, Calif., Jan. 28–30). ACM, New York, 1991, pp. 89–99.

The research of J. R. Driscoll was supported in part by the National Science Foundation (NSF) under grant CCR 88-09573 and by DARPA as monitored by the Air Force Office of Scientific Research under grant AFOSR-90-0292.

The research of D. D. K. Sleator was supported in part by the NSF under grant CCR 86-58139 and by DIMACS (Center for Discrete Mathematics and Theoretical Computer Science), a National Science Foundation Science and Technology Center—NSF-STC 88-09648.

The research of R. E. Tarjan at Princeton University was partially supported by the NSF under grant CCR 89-20505, by the Office of Naval Research (ONR) under grant N00014-91-J-1463, and by DIMACS.

Authors' addresses: J. H. Driscoll, Department of Mathematics and Computer Science, Dartmouth College, Hanover, NH 03755; D. D. K. Sleator, School of Computer Science, Carnegie-Mellon University, Pittsburgh, PA 15213; R. E. Tarjan, Computer Science Department, Princeton University, Princeton, NJ 08544, and NEC Research Institute, Princeton, NJ 08540.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1994 ACM 0004-5411/94/0900-0943 \$03.50

1. Introduction

In this paper we consider the problem of efficiently implementing a set of side-effect-free procedures for manipulating lists. These procedures are:

- makelist*(d): Create and return a new list of length 1 whose first and only element is d .
- first*(X): Return the first element of list X .
- pop*(X): Return a new list that is the same as list X with its first element deleted. This operation does not effect X .
- catenate*(X, Y): Return a new list that is the result of catenating list X and list Y , with X first, followed by Y (X and Y may be the same list). This operation has no effect on X or Y .

We show how to implement these operations so that *makelist*(d) runs in constant time and consumes constant space, *first*(X) runs in constant time and consumes no space, *pop*(X) runs in constant time and consumes constant space, and *catenate*(X, Y) runs in time (and consumes space) proportional to $\log \log k$, where k is the number of list operations done up to the time of the catenation. All of these time and space bounds are amortized over the sequence of operations.

Note. Another important operation on lists is *push*(d, X), which returns the list formed by adding element d to the front of list X . We treat this operation as a special case of catenation, since *catenate*(*makelist*(d), X) has the same effect as *push*(d, X).

This list manipulation problem arises naturally in list-based programming languages such as lisp or scheme. Incorporating our representation for lists into such a language would allow a programmer to use these conceptually simple procedures for manipulating lists without incurring a severe cost penalty. Another application is the implementation of continuations in a functional programming language [Felleisen et al. 1988].

In addition to solving these practical problems very efficiently, our result is the first nontrivial example of a fully persistent data structure that supports an operation that combines two versions. To put this work in context, we need to summarize previous work and terminology on persistence.

A typical data structure allows two types of operations: queries and updates. A query merely retrieves information, whereas an update changes the information represented. Such a data structure is said to be *ephemeral* if queries and updates may only be done on the current version. With an ephemeral structure, the sequence of versions (one for each update done) leading to the current version is lost.

In the *partially persistent* form of an ephemeral data structure, queries may be performed on any version that ever existed, while updates may be performed only on the most recent version. Thus, the collection of versions forms a linear sequence, each one (except the current one) being the predecessor of exactly one version. In the *fully persistent* form of an ephemeral data structure both queries and updates are allowed on any version of the structure. Therefore, one version may be the predecessor of several versions (one for each update that was applied to it). The graph of relations between versions is a tree.

Several researchers have constructed methods for making specific data structures either partially or fully persistent and some general techniques have been developed.¹ In particular, our previous joint work with Neil Sarnak [Driscoll et al. 1989] gives efficient methods for transforming a pointer-based ephemeral data structure into one that is partially or fully persistent in a manner that satisfies ideal resource bounds: a constant factor in query time over that in the ephemeral structure, and a constant amount of space per change in the ephemeral structure. These techniques have several important limitations, however. One of these is that the updates in the ephemeral structure operate on a single version. Combining two versions together, such as catenating two lists, is not allowed. Another limitation is that the in-degree (number of pointers pointing to a particular node) in the ephemeral structure must be bounded by a fixed constant.

A fully persistent version of an ephemeral data structure that supports an update in which two different versions are combined is said to be *confluently persistent*. The relationship between versions of a confluently persistent data structure is that of a directed acyclic graph. The result of this paper, efficient confluently persistent catenable lists, expands the range of data structures to which persistence can be applied.

As a starting point for thinking about this problem, we mention two simple approaches. If we represent the lists as singly-linked lists in the standard way, with a special *entry node* for each version that points to the first element of the version, then two lists X and Y can be catenated in time and space proportional to the length of X . This is done by duplicating the list X and calling the duplicate X' , creating a new entry node that points to X' , and making the last element of X' point to the beginning of Y . Popping the first element of a list is accomplished by creating a new entry node and making it point to the second element of the given list. This takes constant time and space.

This solution is an example of a more general technique known as *path-copying* [Reps et al. 1983]: When a node is to be changed, replace it by a new node, and change all nodes that point to it by applying the same idea recursively. (Such a method will obviously work best when the in-degree of the nodes is small and the paths in the structure are short.) By representing lists as balanced search trees and using the path-copying technique, pop and catenate can be made to run in $O(\log n)$ time and space (n is the length of the longest list involved in the operation). This improves the performance of catenate, but makes popping worse.

In our search for a more efficient solution, we began by trying to apply the *node-copying* method of Driscoll et al. [1989], which can make lists fully persistent using constant time and space per change. This idea breaks down because the method does not allow a pointer in a particular version to be changed to point to a node that is not part of the same version. This is because the node-copying method assigns a *version number* to each version, and maintains a data structure that contains version numbers throughout; to navigate (follow pointers) through a particular version requires comparing the

¹ For examples of partially or fully persistent data structures, see Chazelle [1985], Cole [1986], Dobkin and Munro [1985], Hood and Melville [1981], Krijnen and Meertens [1983], Myers [1982; 1983; 1984], Reps et al. [1983], and Swart [1985]. For examples of general techniques, see Driscoll et al. [1987], Dietz [1989], Overmars [1981], and Sarnak [1986].

version number of that version with those encountered in the traversal, and making decisions based on these comparisons. To accommodate the changing of a pointer to point into a different version would require a new navigation algorithm.

Rather than inventing a more powerful navigation scheme (which does not seem plausible), we instead constructed an efficient list representation that is itself made up of fully persistent (but not efficiently catenatable) lists. These lists are linked together to form a tree whose leaves are the elements of the list to be represented. We devised two solutions based on this underlying representation.

In our first method the subtrees of the root increase in size exponentially from left to right, and these subtrees are *3-collapsible* (a structural property to be defined later that roughly says that the leaves near the left are near the root). We call this the *increasing subtree* method. Because the subtrees increase in size exponentially, only $O(\log n)$ of them are required to represent a list of length n . Popping operates on the leftmost subtree of the root, and makes use of the 3-collapsibility property to run in constant time. Catenating moves the subtrees of the right list into the data structure of the left list one at a time in such a way as to preserve 3-collapsibility and the exponential growth of the subtrees.

Because catenatable lists can get quite large with very few operations (repeatedly catenating a singleton list to itself k times results in a list with length 2^k) we have devised a method that allows this $O(\log n)$ bound to be improved to $O(\log k)$, where k is the total number of list operations.

In our more sophisticated solution (called the *finger tree method*), we change the representation of the list of exponentially growing subtrees, replacing it by a finger search tree. This tree has $O(\log n)$ leaves, so its depth is $O(\log \log n)$. This structure allows popping in constant time and space and catenating two lists of total length n in $O(\log \log n)$ time and space. As before, the n in this bound can be replaced by k , where k is the total number of list operations. This is a dramatic improvement over the naive method, which can require $\Omega(2^k)$ time, and even over an $O(\log n)$ -time catenation method, which can require $\Omega(k)$ time and space per catenate.

This paper is organized as follows: Section 2 outlines the fully persistent memory architecture we shall use, describes the basic representation of lists, and explains the fundamental operations of *delete*, *link*, and *pull* that will be applied to this representation. Section 3 describes the increasing subtree method and the trick for improving the time and space bound for catenation to $O(\log k)$. Section 4 outlines the finger tree method, and Section 5 concludes with open problems and directions for future research.

2. Basic Representation

Our representation of a list is a rooted ordered tree with the list elements in the leaves, such that a left-to-right traversal of the tree visits the elements in the order they occur in the list. Each of the internal nodes of the tree has two or more children. No further structural constraint is imposed on these trees (except implicitly by the operations that manipulate them).

Each internal node of the tree is represented by a header node that has pointers to the first element of a singly-linked list of pointers to its children.

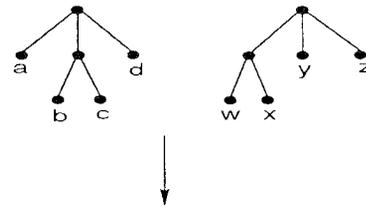


FIG. 2. Linking two trees.

of the root to point to the new child (if this new child is not a leaf), and updating the last three fields of the header node. When a root node becomes a nonroot node, certain information stored there is no longer needed (such as the pointer to the rightmost child, and the links among nonnull child pointers). No harm is done by this extra information. All of these changes can be accomplished in constant time and space; and, because the lists are persistent, the original trees are still available. (See Figure 2.) It is also possible to link the left tree below the root of the right tree, making the root of the left tree the first child of the root of the right tree. This form of linking is useful in the balanced linking method described in Section 3.

It is important to note that the persistence mechanism used to represent the lists does not allow the root list of the first tree to point directly to the beginning of that of the second. If this were allowed, it would solve the problem of efficient catenation. Instead, we must settle for the less satisfactory method described above and work to ensure that the way we link lists together keeps the beginning of the list easily accessible.

The delete operation removes the first child of the root, provided it is a leaf. It is invalid to apply delete to a tree whose first child is not a leaf, and in what follows we shall maintain the invariant that the first child of the root is always a leaf. The operation of deletion can be accomplished in constant time and space while maintaining persistence.

The objective of the pull operation is to bring subtrees up toward the root. It is the tool we shall use to maintain the invariant that the first child of the root is a leaf. Pull takes the first nonleaf child of the root, and “pulls” its leftmost child into the root list. (See Figure 3.) If this creates an internal node with only one child, this child is pulled up as well, and the now childless node is removed from the root list. This is shown in the second pull of Figure 3. If there is no nonleaf child of the root, the tree is said to *flat*. The pull operation applied to a flat tree does nothing. The pointer to the first nonnull child (and the links among the nonnull children) allow this operation to be accomplished in constant time, and to be made fully persistent with only constant space consumption.

LEMMA 1. *The number of pulls required to flatten a tree is the number of leaves minus the degree of the root.*

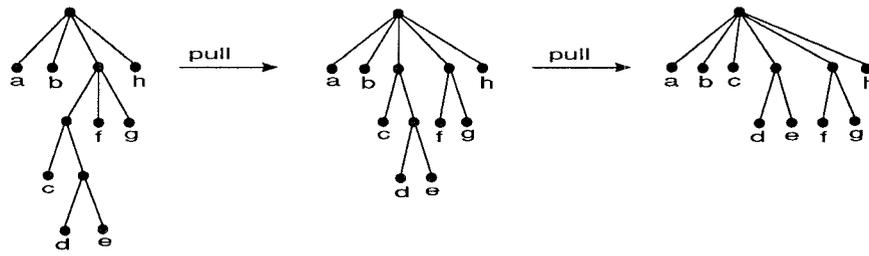


FIG. 3. The pull operation.

This lemma is a consequence of the observation that each pull increases the degree of the root by exactly one.

Note. In our basic representation any change in the tree requires changing all the nodes on the path leading from the location of the change to the root. Therefore, in designing our data structure, we must be wary of making changes deep in the tree.

3. Preliminary Methods

The basic operations can be combined to give several different algorithms with varying efficiency trade-offs. We first describe the *balanced linking* method. At the root of each tree, we keep the size of the tree, which is the number of leaves it contains. We denote the size of a tree T by $|T|$. Catenation is implemented by linking the smaller tree below the larger tree. The pop operation repeatedly pulls the tree until the left child of the root is a leaf. This leaf is then deleted. Catenation takes constant time and space, and pop takes time and space at most proportional to the number of catenations done before the pop.

This algorithm is very closely related to the method of path compression with weighted union for maintaining disjoint sets [Tarjan 1983b]. In fact, it follows from the analysis of path compression that in an ephemeral version of the structure the amortized cost of a pop is $O(\alpha(k))$, where k is the number of list operations. Unfortunately, an amortized bound of this form is simply not sufficient to give efficient full persistence. This is because a “bad” version, one in which pop is expensive, can generate arbitrarily many other bad versions. This would happen if, for example, several different short lists are catenated to the same bad version. Popping each of these new lists will be expensive and will ruin any global amortized bound. There are two ways to circumvent this problem. One is to find some way to make popping globally improve the representations of the lists (rather than locally improving a single version of a list). Another is to find a way to make the amortized bound worst-case. We pursue the latter avenue.

Call a tree *c-collapsible* if repeatedly applying “delete, then pull c times” until the tree becomes empty preserves the invariant that the first child of the root is a leaf. Our goal is to arrange that all trees are *c-collapsible* for some constant c . We now show that *c-collapsibility* is preserved under linking with a sufficiently small tree.

LEMMA 2. *Let T_1 and T_2 be two trees, and let T be the result of linking T_2 to the right of T_1 , making the root of T_2 the last child of the root of T_1 . If T_1 is c -collapsible and $|T_2| \leq (c - 1) \cdot |T_1|$ then T is c -collapsible.*

PROOF. Because T_1 is c -collapsible, $|T_1|$ collapsing steps (each one is a delete followed by c pulls) will work on T (the presence of T_2 does not interfere with the deletions). At this point, $c|T_1|$ pulls have been applied to T . But $|T| = |T_1| + |T_2| \leq c|T_1|$, so $c|T_1|$ pulls are sufficient to flatten T . Therefore, after $|T_1|$ collapsing steps, T is flat; that is, all the leaves are children of the root. \square

In particular, if we choose $c = 3$, then a tree of size up to $2n$ can be linked to a tree of size n while preserving 3-collapsibility. We now use this observation to provide a solution with $O(\log n)$ catenation time.

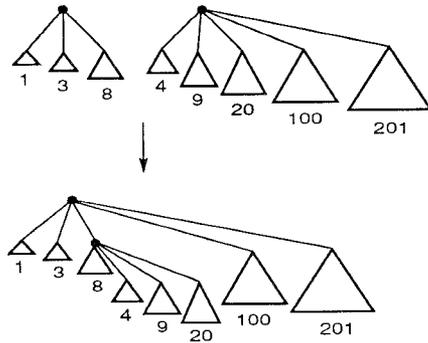
Represent a list by a fully persistent list of trees satisfying the invariant that every tree is 3-collapsible and the i th tree has more than twice as many leaves as the $(i - 1)$ st tree. With each tree in the list, store the size of the tree. The doubling property implies that a list with $t \geq 2$ trees in its representation must have at least 2^t elements. To pop the first element of such a list, apply delete, followed by 3 pulls to the first tree. If this leaves an empty tree, then delete the empty tree from the front of the list of trees. This operation uses constant time and space, and the resulting list has the doubling property.

Catenation is more complicated. To catenate two lists of trees, process each tree of the right list in order from left to right. For each such tree T , compare the size of T with the current size of the rightmost tree T' of the left list. If $|T| \leq 2 \cdot |T'|$, then link T onto T' . Otherwise, make T the new rightmost tree of the left list. (After this happens, no further links will occur in the catenation.) (See Figure 4.) This process preserves the doubling property and 3-collapsibility. Catenation takes $O(\log n)$ time and space because the number of trees representing a list of n elements is at most $\log_2(n + 1)$.

Although catenation has a cost that is logarithmic in the size of the lists involved, a sequence of k catenates can generate a list of length 2^k . Thus, in terms of the number of operations k , the previous result only gives us an $O(k^2)$ bound on the total time and space used to process the sequence.

We can do better in an amortized sense with the following trick: Make an initial guess $k_0 = 2$ of the total number of list operations to be performed. Carry out the operations normally, but keep track of the total size of all trees comprising a list. Keep a record of all the list operations performed as well. If a list becomes longer than $\log_2 k_0$ during a catenation then truncate it so that the length of the resulting list of trees is exactly $\log_2 k_0$. This ensures that catenation will cost $O(\log k_0)$ time and space, and that at least the first k_0 elements of each list will be correctly represented. As long as no more than k_0 pops occur, all the pops will return the correct values. Should the actual number of operations exceed the guess k_0 , then throw away the lists, pick a new guess $k_1 = k_0^2$, and rebuild the lists using the recorded sequence of operations and this new guess.

The following analysis shows that with this method the amortized cost of a pop is $O(1)$ and the total amortized cost of a catenate is $O(\log k)$. Since the guess k_i at least doubles at every phase, the total number of pops, each costing $O(1)$, performed in all phases, is at most $k + k + k/2 + k/4 + \dots = O(k)$. This $O(k)$ cost is a cost of $O(1)$ per operation, whether a pop or a catenate.

FIG. 4. Catenating two lists in $O(\log n)$ time.

Since the final guess k_f is at most k^2 , since $\log_2 k_{i+1} = 2 \log_2 k_i$, and since in each phase the cost of a catenate is $O(\log k_i)$, the total cost of repeating a catenate during each phase is

$$O\left(\sum_{i=0}^f \log k_i\right) = O\left(2 \log k + \log k + \frac{1}{2} \log k + \dots\right) = O(\log k).$$

The same trick can be used to transform the $O(\log n)$ pop and catenate path-copying algorithm mentioned in the introduction into one with cost $O(\log k)$ per operation. This will work even if the length bound is doubled ($k_{i+1} = 2k_i$) rather than being squared at each iteration.

4. The Finger Tree Method

The idea of the finger tree method is to modify the list-of-trees method of Section 3 by replacing the list of collapsible trees² by a balanced binary tree. Each leaf of this balanced tree is the root of one of the collapsible trees. We retain the invariant that each successive collapsible tree is more than twice the size of its predecessor. In the list-of-trees method, the catenation operation linked one collapsible tree at a time until the next one was too big, and then catenated the remainder of the $O(\log n)$ -length list of collapsible trees by copying the list. The point in the list of collapsible trees at which the algorithm stops linking and starts catenating is called the *split point*. The finger tree method uses the same split point, but takes advantage of the balanced tree structure to streamline the process of transferring the collapsible trees.

The algorithm to catenate a structure B to the right of a structure A works roughly as follows: Find the split point in B . Split the balanced binary tree of B in two at the split point, creating B_l and B_r . Link B_l onto the rightmost collapsible tree of A . Join the binary trees of A and B_r together. All of the steps will be done in a fully persistent fashion, and they all will be shown to take time and space proportional to the depth of the balanced binary tree, which is $O(\log \log n)$.

4.1 RED-BLACK TREES. A *red-black tree* [Guibas and Sedgwick 1980; Tarjan 1983b] is a binary tree in which each internal node³ has pointers to both

² For brevity we suppress the “ c ” of “ c -collapsibility” when we wish to avoid specifying a constant. We shall show later that all the trees that we call collapsible are actually 4-collapsible.

³ The *internal nodes* of a binary tree have two children, and the *external nodes* or *leaves* of the tree have no children. Every node of a binary tree is either an internal node or a leaf.

of its children⁴ and a one bit field storing the *color* of the node, which is either red or black.

The colors of the nodes satisfy the following three conditions: (a) Every leaf is black; (b) Every red node has a black parent (or is the root); (c) Every path from the root to a leaf contains the same number of black nodes. For purposes of analyzing red–black trees, it is convenient to assign to each node an integral *rank* that is a function of the coloring of the tree. The rank of each leaf is 0, and the rank of an internal node is equal to the rank of its red children (if any) and one greater than that of its black children (if any). We can interpret the constraints on the coloring of the tree with respect to the ranks of the nodes as follows: (1) The rank of each leaf is 0; (2) The rank of each node that is the parent of a leaf is 1; (3) In traversing from a node to its parent, the rank increases by 0 or 1; (4) In traversing from a node to its grandparent, the rank increases by 1 or 2.

The *depth* of a red–black tree is the maximum number of internal nodes on a path from the root to a leaf. The following lemma (which is a slightly modified form of Lemma 4.1 of Tarjan’s monograph [Tarjan 1983b]) shows that the depth of a red–black tree with n leaves is $O(\log n)$, and will be useful to us later for other purposes.

LEMMA 3. *The subtree rooted at a node of rank r in a red–black tree contains at least 2^r leaves.*

PROOF. If $r = 0$, then the node is a leaf, and the lemma holds. If $r > 0$, then the node is internal and has two children of rank at least $r - 1$. By induction, each of these is the root of a subtree with at least 2^{r-1} leaves. \square

The fundamental operations that we shall need to apply to these trees are inserting a new leaf, deleting a leaf, splitting a tree into two, and joining two trees together. The split operation takes as input a red–black tree along with a particular leaf l , and produces two red–black trees, one containing all the leaves to the left of l , and one containing l and all the leaves to the right of l . The join operation is the inverse of split. It produces a new red–black tree whose leaves are those of the first input tree, followed by those of the second input tree. Algorithms exist to perform all of these operations in time $O(\log n)$ (where n is the number of leaves in all of the trees involved in the operation).

The bottom-up algorithms for insertion and deletion [Tarjan 1983a; 1983b] can roughly be described as follows: A constant amount of restructuring occurs in the vicinity of the change, then color changes are made on some initial part of the path from the change to the root of the tree, then up to three rotations⁵ are done at the top of this path of color changes.

A *finger search tree* is a data structure for representing a list of elements—one of which is a special element called the *finger*—that allows very efficient accesses, insertions and deletions in the vicinity of the finger.⁶ In particular, these three operations can be performed in time $O(\log d)$, where d is the

⁴ Parent pointers may also be included if needed. They are not needed in our application.

⁵ A rotation is a local restructuring operation in a binary tree that maintains the symmetric order of the leaves and changes the depths of various nodes.

⁶ For example, see Brown and Tarjan [1980], Guibas et al. [1977], Kosaraju [1981], and Tsakalidis [1984].

distance in the list between the finger and the location of the operation. The method of Tsakalidis [1984] is particularly useful for our purposes.

In Tsakalidis's method, each element of the list is a leaf in a red–black tree. We assume throughout this discussion that the finger is the leftmost leaf of the tree.⁷ The leftmost path of internal nodes of the tree (along with the internal nodes adjacent to this path) are represented in a special way, but the rest of the tree is just as it would be in a red–black tree.

Tsakalidis's representation of the leftmost path and adjacent internal nodes is called the *spine*. We make use of this data structure without explicitly describing it here. It has the following important properties: (1) The spine consists of nodes with a bounded number of fields; (2) Each node of the spine (except for one entry node) is pointed to only by nodes of the spine; (3) The indegree of any node (the number of pointers to it) is bounded by a constant (independent of the size of the tree); (4) An insertion or deletion operation in the finger search tree causes only a constant number of pointers and fields in the spine to change. This is a consequence of an implicit representation of the colors, and of the $O(1)$ restructuring insertion and deletion algorithms in red–black trees.

These four properties allow us to make this data structure fully persistent by the node-copying method [Driscoll et al. 1989] using only constant space per insertion and deletion, while preserving the $O(\log d)$ running time for these operations.

4.2 THE REPRESENTATION. We are now ready to describe our representation of lists. The list is partitioned into consecutive blocks of elements of sizes b_1, b_2, \dots, b_k . These sizes are arbitrary, subject to the constraint that $b_{i+1} > 2b_i$ for all $1 \leq i < k$. Each block of elements is stored in the leaves of a 4-collapsible tree. Each of these collapsible trees is stored according to the basic representation described in Section 2, and the root of each tree stores its size. A persistent version of Tsakalidis's structure—which we shall call the finger tree—connects the collapsible trees together. The roots of the collapsible trees are the leaves of the finger tree. The internal nodes (except those in the spine data structure) of the finger tree are represented as standard persistent memory elements. As usual, a pointer from one of these nodes to another (or to the root of a collapsible tree) is actually a version number. The spine is stored as a fully persistent version of Tsakalidis' ephemeral representation of the spine.

To find the split point efficiently we store additional information in the internal nodes of the finger tree (except those in the spine). In such a node x , whose descendants are collapsible trees of sizes b_l, b_{l+1}, \dots, b_r , we store the following information: $S_x = \sum_{l \leq i \leq r} b_i$ (the number of leaves that are descendants of this node) and $Q_x = \max_{l \leq i \leq r} \{b_i - 2 \sum_{l \leq m < i} b_m\}$. These quantities can be maintained efficiently under any local restructuring of the tree since, if l and r are the left and right child, respectively, of a node x , we have $S_x = S_l + S_r$ and $Q_x = \max\{Q_l, Q_r - 2S_l\}$.

This completes the static description of our data structure. It remains to describe and analyze pop and catenate.

⁷ In our application we only need this special case. Furthermore, such a solution can be adapted to allow an arbitrary finger location by maintaining two such trees.

4.3 POP. We first observe that the leftmost collapsible tree can be deleted from the structure (or a new leftmost collapsible tree can be inserted into the structure) persistently in constant time and space. This is because the operation on the finger tree is taking place within a constant distance of the finger (0 or 1 for deletion and insertion, respectively), and because the values of S_x and Q_x can be updated in constant time. (Since we do not store these values for nodes in the spine, no updating of these is needed. It is possible, however, that because of the rotations that occur we may introduce a new node outside of the spine for which we must compute these values. This can be done using the update rule above. There are only $O(1)$ such nodes.)

To perform pop, delete the leftmost collapsible tree from the finger tree. Delete the leftmost leaf from it, and apply four pull operations. If the resulting tree is empty, the operation is complete. If it is not empty, then insert the collapsible tree back into the finger tree. All of these steps take constant time and use constant space.

The resulting structure still satisfies all of the required structural constraints; decreasing b_1 maintains $b_2 > 2b_1$, and deleting one element from a 4-collapsible tree and applying four pulls leaves a 4-collapsible tree.

4.4 CATENATE. We describe a catenation method that has the property that the time (and space) required is proportional to the depth of the finger trees involved (which is $O(\log \log n)$ for forming a list of length n), and that maintains all the structural constraints above.

The first step of the algorithm is to transform the two finger trees into red-black trees, in which each internal node is represented by a standard persistent memory element and is endowed with a color and values for S_x and Q_x . This transformation can be done in time and space proportional to the size of the spine of a finger tree, which is proportional to the depth of the finger tree. The colors can be extracted from their implicit representation in the spine, and the data values can be computed from those of their children, which are available.

Let the two structures to be catenated be A and B , and let the collapsible trees of A be A_1, A_2, \dots, A_f of sizes a_1, a_2, \dots, a_f , and the collapsible trees of B be B_1, B_2, \dots, B_k of sizes b_1, b_2, \dots, b_k . The split point is defined to be the least $i \geq 1$ such that $b_i > 2(a_f + \sum_{1 \leq j \leq i-1} b_j)$ (or equivalently $2a_f < b_i - 2 \sum_{1 \leq j \leq i-1} b_j$). If this inequality is not true for any $i \leq k$, then the split point is defined to be $k + 1$. We can now describe how to determine if the split point is in a subtree T rooted at the node x given S_x (the total size of the trees to the left of the subtree T) and a_f . If $2a_f < Q_x - 2S_x$, then a point i satisfying the above inequalities is in T ; otherwise, there is no such point in T . Using this test, it is easy to locate the split point in time proportional to the depth of the red-black tree: merely walk down from the root (updating S_x), always choosing the leftmost option that contains a point satisfying the inequality. The leaf we reach in this way is the split point.

Once the split point i has been found, we split the red-black tree at this point, creating a new red-black tree T_l with collapsible trees B_1, B_2, \dots, B_{i-1} , and a new red-black tree T_r with collapsible trees B_i, B_{i+1}, \dots, B_k . (If the split point is $k + 1$, then T_l is empty. If the split point is 1, then T_r is empty.) We now change the representation of the leftmost path of T_l into a single persistent memory element, link the resulting structure to the right of A_f , and

call the result A'_f . (The change in representation is required to guarantee the collapsibility of A'_f .) The link forming A'_f occurs deep in the red–black tree of A . Making such a change persistent requires that all the nodes on the path from A_f to the root be changed (just as in path-copying). Furthermore, the S_x and Q_x fields of all the nodes on this path must be updated. The cost of these changes is proportional to the depth of the red–black tree. Finally, we join the new version of the red–black tree of A with the red–black tree of T_r , and transform the resulting red–black tree into a finger tree (by building the spine).⁸ The process is shown schematically in Figure 5.

The $O(\log \log n)$ bound on the time and space cost of catenate follows immediately from the structural constraints on these trees. It remains to prove that these constraints are preserved by our algorithm.

In the data structure resulting from a catenation, each collapsible tree is more than twice the size of its predecessor. To verify this, we need only check the two places where this constraint might be violated. These are between A_{f-1} and A'_f and between A'_f and B_i . The former place satisfies the constraint since $a'_f \geq a_f$. By our choice of the split point, we know that $b_i > 2(a_f + b_1 + \dots + b_{i-1}) = 2a'_f$, which proves that the latter point also satisfies the constraint.

The only remaining detail is to show that the tree A'_f is 4-collapsible. This requires a more detailed analysis of the structure of red–black trees and properties of the pull operation.

LEMMA 4. *Let l be the i th leaf ($i \geq 2$) of a red–black tree. The distance between l and the nearest internal node on the left path is at most $2\lceil \log i \rceil$.*

PROOF. Let $r = \lceil \log i \rceil$. Let x be the deepest node on the left path of the red–black tree that is of rank r . By Lemma 3, the subtree rooted at x contains at least 2^r leaves. These leaves are consecutive starting from the leftmost leaf. Since $2^r \geq i$, this subtree must contain l , so x is an ancestor of l .

In traversing the path from l toward the root of the tree, the first step goes from a node of rank 0 to one of rank 1. Subsequently, every pair of steps increases the rank by at least 1. Therefore, after $2r + 1$ steps, a node of rank at least $r + 1$ must be reached. Since x is on this path and has rank r , it must have been passed. Therefore, x is reached after at most $2r$ steps. \square

LEMMA 5. *For any j ($1 \leq j \leq i - 1$), let l be the number of leaves to the left of B_j in A'_f . The number of pulls that must be applied to A'_f until a state is reached in which B_j and all the leaves to the left of B_j are adjacent to the root of the tree is at most $2l$.*

PROOF. Let F be the tree (reached by applying pulls to A'_f) in which B_j and all of the l leaves to the left of B_j are adjacent to the root. (See Figure 6.)

⁸ In linking T_i into A_f , we have in a sense mixed apples and oranges, since the memory elements of the nodes of T_i contain color and size fields. After the link is done, the data in these fields is no longer needed. No harm is done by the presence of this extra useless information in the tree. A slightly different approach is to completely avoid the use of the persistent memory architecture in the nodes of the red–black tree. This approach works, but a different non-uniformity results when the link is done. Now the resulting tree contains a mixture of nodes, some of which are ordinary red–black tree nodes, while others are persistent nodes. It is too expensive to clean up the entire tree immediately. Instead the pop algorithm must replace the red–black nodes by persistent nodes gradually as they are encountered.

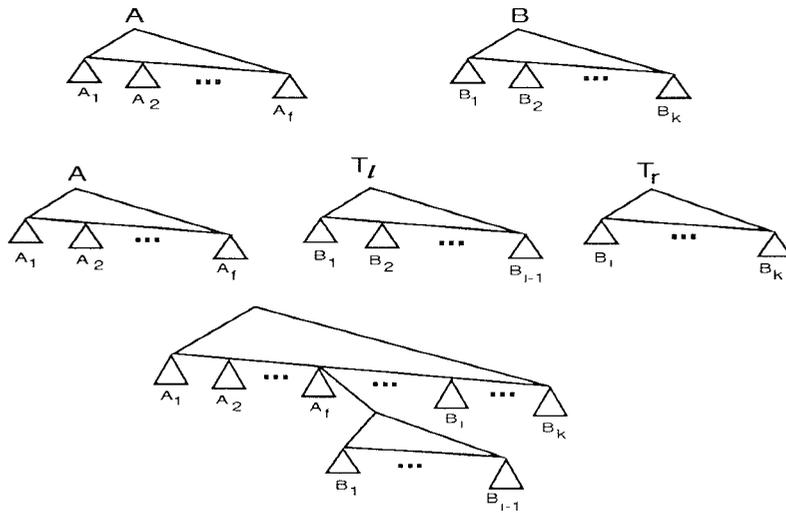


FIG. 5. The top picture above shows the initial trees A and B . The center picture is an intermediate point in concatenating A and B . The bottom picture shows the final result.

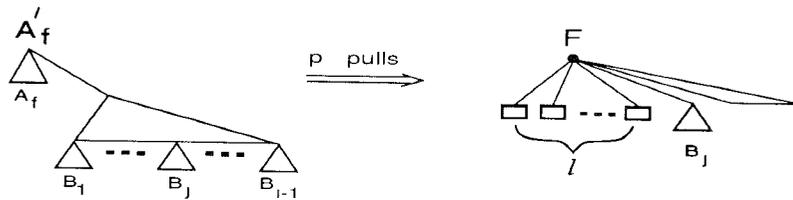


FIG. 6. The process of transforming A'_f to F by applying p pulls.

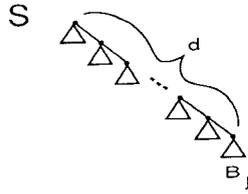
Let p be the number of pulls required to reach this state starting from tree A'_f . Our goal is to bound p by $2l$.

The *stripped tree* S is obtained by starting with A'_f and deleting all nodes that are not ancestors of a leaf in B_j or a leaf to the left of B_j . The path from the root of B_j to the root of S is a rightmost path of S . (See Figure 7.) Let d be the distance between the root of B_j and the root of S (or equivalently the root of A'_f). Let q be the number of nodes on this path that have only one child. (Although all internal nodes in A'_f have at least two children, this may not be true in S .)

Let F' be the stripped form of F . We shall now evaluate p' , the number of pulls required to transform S into F' . Each pull either increases the degree of the root by one or deletes an internal node with one child. Therefore,

$$\begin{aligned}
 p' &= (\text{degree of root of } F') - (\text{degree of root of } S) + q \\
 &= 1 + l + q - (\text{degree of root of } S) \\
 &\leq 1 + l + q - 2 = l + q - 1.
 \end{aligned}$$

What is the relationship between p and p' ? Stripping accelerates the pulling process by at most one for every node on the path from the parent of the root of B_j to the root of S , except for the root and the nodes on this path with one

FIG. 7. The stripped tree S .

child. This is because for the nodes with more than one child, the pull that promotes the subtree immediately to the left of the path also promotes the subtree containing B_j . This extra promotion may not happen in the pull sequence in the unstripped tree. (It depends on whether stripping removed a subtree to the right of the node or not.) To summarize, we have

$$p \leq p' + d - 1 - q \leq l + d - 2.$$

Now we observe that $d \leq j + 2$. Lemma 4 shows that the distance between the root of B_j and the root of T_l (after the leftmost path of T_l is made into a single memory element) is at most $2\lceil \log j \rceil$. It is easy to verify that $2\lceil \log j \rceil \leq j + 1$ for all $j \geq 2$. For $j = 1$, the distance between the root of B_1 and the root of A'_f is 2. This completes the proof of the observation.

Also note that $j \leq l$, since each of $A_f, B_1, B_2, \dots, B_{j-1}$ has at least one leaf. Combining these observations gives

$$p \leq l + d - 2 \leq l + j + 2 - 2 \leq l + j \leq 2l. \quad \square$$

LEMMA 6. *The tree A'_f is 4-collapsible.*

PROOF. The number of pulls required to transform F (of the previous lemma) into a tree in which all the leaves of B_j are adjacent to the root is at most b_j . By the choice of the split point, we know that $b_j \leq 2l$.

When repeatedly applying pop to the tree A'_f , by the time the first l leaves have been deleted, $4l$ pulls have been done. By Lemma 5, the first $2l$ of these pulls suffice to create the tree F , and the remaining ones suffice to collapse B_j . Thus, by the time any leaves of B_j need to be popped, all of the leaves of B_j are adjacent to the root.

We now consider the entire process of popping A'_f . The first a_f pops work because the tree A_f is 4-collapsible. It follows from the discussion above that by the time all of the leaves of A_f have been popped, all of the leaves of B_1 are adjacent to the root. By the time these leaves have been deleted, all the leaves of B_2 are adjacent to the root. This argument applies for all B_j , so we conclude that A'_f is 4-collapsible. \square

As before, we can convert the $O(\log \log n)$ catenation time to amortized $O(\log \log k)$. In the previous version, each k_i was squared, thereby doubling $\log_2 k_i$. Here, we choose $k_i = 2^{2^{2^i}}$, thereby doubling $\log \log k_i$ at each new guess. Truncation can be used to ensure that the height of the finger tree is $O(\log \log k_i)$. (From a practical perspective, an additional detail to be considered is that the size of the guesses grows so rapidly that some guess may become too large to represent and manipulate efficiently in a real computer. If this happens, just truncate the lists at the maximum representable size. In such a setting, it is unlikely that the number of operations would exceed the size of the largest representable number.)

5. Remarks

We have shown how to make lists persistent under the operations of pop and catenate, where any two versions of a list may be catenated. Although the time bounds for the operations are low, we have left a variety of interesting open problems to explore.

A large amount of machinery was used to achieve the $O(\log \log k)$ result, which means that it is probably not very useful in practice. Is there a simpler method to achieve the same bound? From a practical perspective the best method to use is perhaps path-copying in a binary search tree, with the doubling trick mentioned at the end of Section 3.

There is no nontrivial lower bound known for this problem. Our intuition is that there may well be a way of catenating in constant time. It is an intriguing open problem to find such a method.

Another direction for future work is to make other data structures confluent persistent. A close relative of our lists is the catenable deque, which allows pops from both ends. We have devised an alternative way of representing catenable lists that extends to the problem of catenable deques, with somewhat worse resource bounds. The idea is to use red-black trees with lazy recoloring [Driscoll et al. 1989] to represent the deques, with certain paths in the trees themselves represented by balanced search trees. The resulting data structure supports pop from either end and catenate in time and space $O(\log \log n)$, or amortized time and space $O(\log \log k)$ using a variant of the list truncation trick. The amortization in this and in all our $O(\log \log k)$ results can be eliminated by using incremental rebuilding of the data structures.

Since the conference version of the present paper appeared, work by Buchsbaum and Tarjan [1993] has shed additional light on the problem of catenable deques. They use a bootstrapping technique to represent a catenable deque by a balanced tree in which certain paths are themselves represented by catenable deques. The best version of their method has an $O(1)$ worst-case time bound for catenate and an $O(\log^*k)$ worst-case time bound for a pop from either end of a deque. By combining their method with ours, one can obtain for the single-ended problem worst-case bounds of $O(1)$ for pop and $O(\log^{(i)}k)$ for catenate for any constant i . Note that our method gives $O(1)$ time for pop, while theirs gives $O(1)$ time for catenate, but obtaining both simultaneously remains an open problem.

A bigger challenge is to devise a general technique for automatically making ephemeral combinable data structures confluent persistent. Such a result would be very useful in the implementation of high-level languages. It would allow, for instance, the efficient passing of large linked data structures as call-by-value parameters since the data structure would not need to be copied, but could be modified by means of persistence at low cost per change.

ACKNOWLEDGMENT. Although the problem addressed in this paper appeared as an open problem in Driscoll et al. [1989], we were motivated to pursue it actively by Mark Krentel, who pointed out its application to continuations.

REFERENCES

- BROWN, M. R., AND TARJAN, R. E. 1980. Design and analysis of a data structure for representing sorted lists. *SIAM J. Comput.* 9 (1980), 594–614.
- BUCHSBAUM, A. L., AND TARJAN, R. E. 1993. Confluent persistent deques via data structural

- boot-strapping. In *Proceedings of the 4th Annual ACM-SIAM Symposium on Discrete Algorithms*. ACM, New York, pp. 155–164.
- CHAZELLE, B. 1985. How to search in history. *Inf. Control* 77, 77–99.
- COLE, R. 1986. Searching and storing similar lists. *J. Algorithms* 7, 202–220.
- DIETZ, P. F. 1989. Fully persistent arrays. *Proceedings of the 1989 Workshop on Algorithms and Data Structures* (Ottawa, Ont., Canada). *Lecture Notes in Computer Science*. Vol. 382. Springer-Verlag, New York, pp. 67–74.
- DOBKIN, D. P., AND MUNRO, J. I. 1985. Efficient uses of the past. *J. Algorithms* 6, 455–465.
- DRISCOLL, J., SARNAK, N., SLEATOR, D., AND TARJAN, R. 1989. Making data structures persistent. *J. Comput. Syst. Sci.* 38, 86–124.
- FELLEISEN, M., WAND, M., FRIEDMAN, D. P., AND DUBA, B. F. 1988. Abstract continuations: A mathematical semantics for handling full functional jumps. In *Proceedings of the 1988 ACM Conference on Lisp and Functional Programming* (Snowbird, UT, July 25–27). ACM, New York, pp. 52–62.
- GUIBAS, L. J., MCCREIGHT, E. M., PLASS, M. F., AND ROBERTS, J. R. 1977. A new representation for linear lists. In *Proceedings of the 9th Annual ACM Symposium on Theory of Computing* (Boulder, Co., May 2–4). ACM, New York, pp. 49–60.
- GUIBAS, L. J., AND SEDGEWICK, R. 1980. A dichromatic framework for balanced trees. In *Proceedings of the 21st Annual IEEE Symposium on the Foundations of Computer Science*. IEEE, New York, pp. 8–21.
- HOOD, R., AND MELVILLE, R. 1981. Real-time queue operations in pure LISP. *Inf. Proc. Lett.* 13, 50–54.
- KOSARAJU, R. 1981. Localized search in sorted lists. In *Proceedings of the 13th Annual ACM Symposium on Theory of Computing* (Milwaukee, Wis., May 11–13). ACM, New York, pp. 62–69.
- KRIJNEN, T., AND MEERTENS, L. G. L. T. 1983. Making B-trees work for B. IW 219 83, The Mathematical Centre. Amsterdam, The Netherlands.
- MYERS, E. W. 1982. AVL dags. Tech. Rep. TR 82-9. Dept. Computer Science, Univ. Arizona, Tucson, Az.
- MYERS, E. W. 1983. An applicative random-access stack. *Inf. Proc. Lett.* 17, 241–248.
- MYERS, E. W. 1984. Efficient applicative data types. In *Conference Record of the 11th Annual ACM Symposium on Principles of Programming Languages* (Salt Lake City, Ut., June 15–18). ACM, New York, pp. 66–75.
- OVERMARS, M. H. 1981. Searching in the past, II: General Transforms. Tech. Rep. RUU-CS-81-9. Department of Computer Science, Univ. Utrecht, Utrecht, The Netherlands.
- REPS, T., TEITELBAUM, T., AND DEMERS, A. 1983. Incremental context-dependent analysis for language based editors. *ACM Trans. Prog. Lang. Systems* 5, 449–477.
- SARNAK, N. 1986. Persistent Data Structures. Ph.D. dissertation, Department of Computer Science, New York Univ., New York.
- SWART, G. F. 1985. Efficient Algorithms for Computing Geometric Intersections. Tech. Rep. 85-01-02. Department of Computer Science, Univ. Washington, Seattle, Wash.
- TARJAN, R. E. 1983a. Updating a balanced search tree in $O(1)$ rotations. *Inf. Proc. Lett.* 16, 253–257.
- TARJAN, R. E. 1983b. *Data structures and network algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, Pa.
- TSAKALIDIS, A. K. 1984. An optimal implementation for localized search. A 84/06, Fachbereich Angewandte Mathematik und Informatik, Universität des Saarlandes, Saarbrücken, West Germany.

RECEIVED JANUARY 1992; REVISED JULY 1993; ACCEPTED JULY 1993