

Balanced Parentheses Strike Back

HSUEH-I LU AND CHIA-CHI YEH

National Taiwan University

Abstract. An *ordinal tree* is an arbitrary rooted tree where the children of each node are ordered. Succinct representations for ordinal trees with efficient query support have been extensively studied. The best previously known result is due to Geary et al. [2004b, pages 1–10]. The number of bits required by their representation for an n -node ordinal tree T is $2n + o(n)$, whose first-order term is information-theoretically optimal. Their representation supports a large set of $O(1)$ -time queries on T . Based upon a balanced string of $2n$ parentheses, we give an improved $2n + o(n)$ -bit representation for T . Our improvement is two-fold: First, the set of $O(1)$ -time queries supported by our representation is a proper superset of that supported by the representation of Geary, Raman, and Raman. Second, it is also much easier for our representation to support new queries by simply adding new auxiliary strings.

Categories and Subject Descriptors: E.1 [Data Structure]: —Trees; E.4 [Coding and Information Theory]: —Data compaction and compression; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—Computations on discrete structures; G.2.2 [Discrete Mathematics]: Graph Theory—Graph algorithms, Trees; H.3.1 [Information Storage and Retrieval]: Content Analysis and Indexing—Dictionaries, indexing methods

General Terms: Algorithm, Design, Theory

Additional Key Words and Phrases: Succinct data structures, XML document representation

ACM Reference Format:

Lu, H.-I., and Yeh, C.-C. 2008. Balanced parentheses strike back. *ACM Trans. Algor.* 4, 3, Article 28, (June 2008), 13 pages. DOI = 10.1145/1367064.1367068 <http://doi.acm.org/10.1145/1367064.1367068>

1. Introduction

An *ordinal tree* (see, e.g., Geary et al. [2004b] and Benoit et al. [2005]) is an arbitrary rooted tree where the children of each node are ordered. All trees in this

This research was supported in part by NSC Grants 94-2213-F-002-126, 95-2221-E-002-077, and 96-2221-E-002-033.

H.-I. Lu is the corresponding author and is also affiliated with the Graduate Institute of Networking and Multimedia and the Graduate Institute of Biomedical Electronics and Bioinformatics, National Taiwan University.

Authors' address: Department of Computer Science and Information Engineering, National Taiwan University, 1 Roosevelt Road, Section 4, Taipei 106, Taiwan, Republic of China, e-mail: {hil; r93048}@csie.ntu.edu.tw.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2008 ACM 1549-6325/2008/06-ART28 \$5.00 DOI 10.1145/1367064.1367068 <http://doi.acm.org/10.1145/1367064.1367068>




FIG. 1. Three representations for the same tree.

article are ordinal. The number of distinct n -node trees is $2^{2n - \Theta(\log n)}$ [Graham et al. 1989], so the information-theoretically minimum number of bits to differentiate these trees is $2n - \Theta(\log n)$. There are three major types of $2n$ -bit representations for an n -node tree T :

- Balanced parentheses [Munro and Raman 2001; Chuang et al. 1998; He et al. 1999; Chiang et al. 2005; Munro and Rao 2004; Bonichon et al. 2006], a folklore encoding consisting of a balanced string of parentheses representing the counterclockwise depth-first traversal of T , where an open (respectively, closed) parenthesis denotes a descending (respectively, ascending) edge traversal. For technical reason, one usually adds a pair of enclosing parentheses to the above $2n - 2$ parentheses, resulting in a representation consisting of $2n$ parentheses.
- Level order unary degree sequence (LOUDS) [Jacobson 1989], representing a node of degree d as a string of d copies of 1-bits followed by a 0-bit, where these nodes are represented in a level-order traversal of T .
- Depth first unary degree sequence (DFUDS) [Benoit et al. 2005], representing a node of degree d as a string of d copies of 1-bits followed by a 0-bit, where these nodes are represented in a depth-first traversal of T .

An example is shown in Figure 1.

Initiated by Jacobson [1989], succinct representations for trees with efficient query support have been extensively studied in the literature. Jacobson [1989] extended the LOUDS representation into a $\Theta(n)$ -bit encoding to support the parent query and the rank and select queries for nodes in level-order traversal of T in $\Theta(\log n)$ time. Clark and Munro [Clark 1996; Clark and Munro 1996] squeezed Jacobson's encoding into a $3n + o(n)$ -bit representation, from which the above queries and the subtree-size query can be supported in $O(1)$ time. Later succinct representations, all have $2n + o(n)$ bits, form the following trade-off between the choices of base representations and the sets of supported $O(1)$ -time queries:

- Based upon balanced parentheses, Munro and Raman [2001] showed that an $o(n)$ -bit auxiliary string suffices to support the following queries in $O(1)$ time: parent, depth, subtree-size, and the rank and select queries for nodes in pre-order and post-order traversal of T . Munro et al. [2001] showed an $o(n)$ -bit auxiliary string to support $O(1)$ -time query for leaf-rank, leaf-select, and leaf-size. Chiang et al. [2005] showed an $o(n)$ -bit auxiliary string to support $O(1)$ -time degree query. Munro and Rao [2004] further gave an $o(n)$ -bit auxiliary string to support $O(1)$ -time level-ancestor query.

TABLE I. A SUMMARY FOR CURRENT $2n + o(n)$ -BIT ENCODINGS FOR AN n -NODE TREE: PARENTHESES [MUNRO AND RAMAN 2001, 2004; CHIANG ET AL. 2005; MUNRO AND RAO 2004; MUNRO ET AL. 2001], DFUDS [BENOIT ET AL. 2005], GEARY ET AL. [2004a]

	Parentheses	DFUDS	Geary et al.	new
pre-order select and rank	✓	✓	✓	✓
post-order select and rank	✓		✓	✓
child-select and child-rank		✓	✓	✓
leaf-select, leaf-rank, and leaf-size	✓			✓
lowest common ancestor				✓
subtree height				✓
subtree size	✓	✓	✓	✓
level ancestor	✓		✓	✓
distance				✓
degree	✓	✓	✓	✓
depth	✓		✓	✓

—Based upon the DFUDS representation, Benoit et al. [2005] gave an $o(n)$ -bit auxiliary string that supports the following queries in $O(1)$ time: child-rank, child-select, degree, subtree-size, and node-rank and node-select in the pre-order traversal of T . However, such a choice of the base representation still does not provide $O(1)$ -time support for the depth and level-ancestor queries, the node-rank and node-select queries in the post-order traversal of T , and the rank, select, and size queries for leaves.

Recently, Geary et al. [2004b] almost resolved the above trade-off by giving a $2n + o(n)$ -bit encoding for T that supports in $O(1)$ time the aforementioned queries except those leaf-related ones [Munro et al. 2001]. Their approach differs from all previous work achieving $2n + o(n)$ bits in that their encoding does not consist of a $2n$ -bit base representation for the topology of T plus an $o(n)$ -bit auxiliary string. Instead, they decomposed T into several types of subtrees, whose topologies are represented in a hierarchical way, where different levels are composed of mixtures of different base representations and auxiliary strings. Such an involved structure seriously complicates the possibility of supporting additional queries using other stand-alone auxiliary strings. An implementation based upon a similar concept is studied in Geary et al. [2004b]. Very recently, Delpratt et al. [2006] showed that LOUDS-based representation can also be implemented to have competitive practical performance.

In this article, we give new $o(n)$ -bit auxiliary strings for the $2n$ -bit balanced string of parentheses representing T . Together with previous $o(n)$ -bit auxiliary strings for balanced parentheses [Munro and Raman 2001; Chiang et al. 2005; Munro and Rao 2004], our $2n + o(n)$ -bit encoding for T supports all of Geary et al.'s queries in $O(1)$ time. Consisting of a base representation plus $o(n)$ -bit auxiliary strings, our encoding is better in the ease of supporting new queries by adding new $o(n)$ -bit auxiliary strings. To demonstrate such an advantage, we also show how to handle $O(1)$ -time queries currently unsupported by Geary et al.'s encoding, including (a) lowest common ancestor, (b) distance, and (c) subtree height. Table I summarizes the above discussion.

We follow the convention of unit-cost RAM model of computation with $\Theta(\log n)$ -bit word size [van Emde Boas 1990], which is assumed in all the previous work except that of Jacobson [1989]. The rest of this article is organized as follows.

Section 2 gives the preliminaries. Section 3 shows our auxiliary strings for distance, subtree height, and lowest common ancestor. Section 4 shows our auxiliary strings for child-rank and child-select.

2. Preliminaries

Let T be the input n -node tree. Let v_i denote the i th node of T in the pre-order traversal of T . Let S be the balanced string of $2n$ parentheses for T . Let $S[i, j]$ denote the substring of S from index i to index j . Let $S[i] = S[i, i]$. Let ℓ_i be the index such that $S[\ell_i]$ is the i -th open parenthesis in S . Let r_i be the index such that $S[r_i]$ is the closed parenthesis that matches $S[\ell_i]$ in S . One can easily see that the correspondence between v_i and the matched parentheses $S[\ell_i]$ and $S[r_i]$: v_i is the parent of v_j if and only if $S[\ell_i]$ and $S[r_i]$ is the closest parenthesis pair that encloses $S[\ell_j]$ and $S[r_j]$. Let $w(i, j) = j - i + 1$. For the rest of the paper, all logarithms are of base 2. Let $B = \lceil \log^3 n \rceil$, $b = \lceil (\log \log n)^3 \rceil$, $n_B = \lceil \frac{2n}{B} \rceil$, and $n_b = \lceil \frac{2n}{b} \rceil$.

LEMMA 2.1 (SEE BELL ET AL. [1990] AND ELIAS [1975]). *For any $O(n)$ -bit strings S_1, S_2, \dots, S_k with $k = O(1)$, there is an $O(\log n)$ -bit auxiliary string α_{concat} such that, given the concatenation of $\alpha_{concat}, S_1, S_2, \dots, S_k$ as input, the index of the first symbol of any given S_i in the concatenation is computable in $O(1)$ time.*

Let $S_1 \circ S_2 \circ \dots \circ S_k$ denote the concatenation of $\alpha_{concat}, S_1, S_2, \dots, S_k$ as in Lemma 2.1.

LEMMA 2.2 (SEE MUNRO AND RAMAN [2001] AND CHIANG ET AL. [2005]). *Let S be a length- $2n$ string of balanced parentheses that represents an n -node tree T . It takes $O(n)$ time to compute an $o(n)$ -bit string α_{aux} such that the following queries for S can be determined from S and α_{aux} in $O(1)$ time: (a) the parent, degree, and depth of v_i in T , (b) the parenthesis that matches $S[i]$ in S , and (c) the rank and select queries for open and closed parentheses in S .*

By Lemma 2.2, given $S \circ \alpha_{aux}$, indices i , ℓ_i , and r_i can be determined from one another in $O(1)$ time. Our technique of dividing the input strings into multiple levels of blocks, which has been widely used in many succinct data structures, is inspired by Munro and Raman [Munro 1996; Munro and Raman 2001].

3. Distance, Subtree Height, and Lowest Common Ancestor

Let L be the $2n$ -element array such that each $L[i]$ is the number of open parentheses minus the number of closed parentheses in $S[1, i]$. Therefore, if $S[j]$ is the i th open parenthesis in S , then $L[j]$ is the level of v_i in T . For any indices i and j with $i \leq j$, let $index_{min}(L, i, j)$ (respectively, $index_{max}(L, i, j)$) denote the smallest index k with $i \leq k \leq j$ such that $L[k]$ equals the minimum (respectively, maximum) of $L[i], L[i+1], \dots, L[j]$. As observed by Gabow et al. [1984], the lowest-common-ancestor query can be reduced to the above range-minima query $index_{min}$. Similarly, our auxiliary string for supporting the queries of distance, subtree height, and lowest common ancestor is based on the lemma below. Observe that each $L[i]$ can be obtained from S in $O(1)$ time using the auxiliary string α_{aux} for the rank queries with respect to open and

closed parentheses in S . Therefore, the following lemma does not require L in the encoding.

Let I be an array of m indices. Let $k_{\min}(I, m, i, j)$ (respectively, $k_{\max}(I, m, i, j)$) be the smallest index k with $i \leq k \leq j$ that minimizes (respectively, maximizes) $L[I[k]]$. We first prove the following lemma using techniques extended from Section 3 of Bender and Farach-Colton [2000].

LEMMA 3.1. *It takes $O(m \log m)$ time to compute an $O(m \log^2 m)$ -bit string $\alpha_q(I, m)$ from which $k_{\min}(I, m, i, j)$ and $k_{\max}(I, m, i, j)$ for any indices i and j with $1 \leq i \leq j \leq m$ can be determined from S , α_{aux} , and α_q in $O(1)$ time.*

PROOF. For each $i = 1, 2, \dots, m$ and $j = 1, 2, \dots, \lceil \log m \rceil$, let $M_{\min}[i][j]$ (respectively, $M_{\max}[i][j]$) be the smallest index k with $i \leq k < i + 2^j$ that minimizes (respectively, maximizes) $L[I[k]]$. Let $\alpha_q(I, m) = M_{\min} \circ M_{\max}$. Observe that $\alpha_q(I, m)$ takes $O(m \log^2 m)$ bits and can be computed from L and I in $O(m \log m)$ time using dynamic programming. Let $k_1 = M_{\min}[i][k]$ and $k_2 = M_{\min}[j - 2^k + 1][k]$, where $k = \lfloor \log(j - i) \rfloor$. It is not difficult to see that

$$k_{\min}(I, m, i, j) = \begin{cases} k_1 & \text{if } L[I[k_1]] < L[I[k_2]] \\ k_2 & \text{otherwise.} \end{cases}$$

One can compute $k_{\max}(I, m, i, j)$ from M_{\max} , I , and L analogously in $O(1)$ time. \square

LEMMA 3.2. *It takes $O(n)$ time to compute an $o(n)$ -bit string α_{rmq} such that $\text{index}_{\min}(L, i, j)$ and $\text{index}_{\max}(L, i, j)$ for any indices i and j can be computed from S , α_{aux} , and α_{rmq} in $O(1)$ time.*

PROOF. First let I_B be the n_B -element array such that each $I_B[i]$ is the smallest index j with $(i-1)B < j \leq iB$ that minimizes $L[j]$. I_B takes $O(n_B \log B) = o(n)$ bits. Also, for each $i = 1, 2, \dots, n_B$, let $I_b[i]$ be the $\lceil \frac{B}{b} \rceil$ -element array such that each $I_b[i][j]$ is the smallest index t with $(j-1)b < t \leq jb$ that minimizes $L[(i-1)B + t]$. I_b takes $O(n_B \lceil \frac{B}{b} \rceil \log b) = o(n)$ bits. Let $\alpha_{q1} = \alpha_q(I_B, n_B)$, and for each $i = 1, 2, \dots, n_B$, let $\alpha_{q2}[i] = \alpha_q(I_b[i], \lceil \frac{B}{b} \rceil)$. By Lemma 3.1, both of α_{q1} and α_{q2} take $o(n)$ bits and can be obtained in $O(n)$ time. Finally, let α_{q3} be an $O(n)$ -time obtainable table such that any $\text{index}_{\min}(L, i, j)$ and $\text{index}_{\max}(L, i, j)$ with $w(i, j) \leq 2b$ can be computed from $S[i, j]$ and α_{q3} in $O(1)$ time. That is, let $\alpha_{q3}[S[i, i+2b-1]][j-i+1] = (\text{index}_{\min}(L, i, j) - i, \text{index}_{\max}(L, i, j) - i)$ for any indices i and j with $w(i, j) \leq 2b$. Since each entry takes $O(\log b)$ bits, the number of bits required by α_{q3} is $O(2^{2b} 2b \log b) = o(n)$. Let $\alpha_{rmq} = \alpha_{q1} \circ \alpha_{q2} \circ \alpha_{q3} \circ I_B \circ I_b$, which has $o(n)$ bits and is obtainable in $O(n)$ time.

To answer $\text{index}_{\min}(L, i, j)$ from S , α_{aux} , and α_{rmq} , we can always decompose the interval $[i, j]$ into two (not necessarily disjoint) subintervals $[i_1, j_1]$ and $[i_2, j_2]$ whose union is $[i, j]$. Clearly $\text{index}_{\min}(L, i, j)$ can be determined from $\text{index}_{\min}(L, i_1, j_1)$ and $\text{index}_{\min}(L, i_2, j_2)$ in $O(1)$ time. Consider the following cases.

- Case 1. $w(i, j) \leq 2b$. We simply resort to $S[i, j]$ and α_{q3} .
- Case 2. $w(i, j) > 2b$ and $S[i, j]$ is in the same length- B block of S . Since $\text{index}_{\min}(L, i, i+b-1)$ and $\text{index}_{\min}(L, j-b+1, j)$ can be determined in $O(1)$ time using Case 1, it suffices to determine $\text{index}_{\min}(L, i', j')$, where (a) i' is the smallest index with $i \leq i'$ that is a starting index of a length- b block of S , and (b) j' is the largest index with $j' \leq j$ that is an ending index of a length- b block

- of S . Since i' and j' are in the same length- B block of S , $\text{index}_{\min}(L, i', j')$ can be determined from S , α_{aux} , and α_{q2} in $O(1)$ time.
- Case 3. $w(i, j) > 2b$ and $S[i, j]$ belongs to two or more consecutive length- B blocks of S . Let $i' - 1$ be the ending index of the length- B block of S that contains i . Let $j' + 1$ be the starting index of the length- B block of S that contains j . Since $\text{index}_{\min}(L, i, i' - 1)$ and $\text{index}_{\min}(L, j' + 1, j)$ can be determined in $O(1)$ time using Case 2, it suffices to determine $\text{index}_{\min}(L, i', j')$ for the case that $i' \leq j'$. Since i' is a starting index of a length- B block of S and j' is an ending index of a length- B block of S , one can determine $\text{index}_{\min}(L, i', j')$ from S , α_{aux} , and α_{q1} in $O(1)$ time.

It is not difficult to answer $\text{index}_{\max}(L, i, j)$ from S , α_{aux} , and α_{rmq} analogously in $O(1)$ time. \square

As pointed out by an anonymous reviewer, our data structure for lowest common ancestor is similar to that of Sadakane [2002] for suffix arrays.

THEOREM 3.3. *It takes $O(n)$ time to compute an $o(n)$ -bit string α_{new1} such that the queries of distance, subtree height, and lowest common ancestor can be answered from S and α_{new1} in $O(1)$ time.*

PROOF. Let $\alpha_{new1} = \alpha_{aux} \circ \alpha_{rmq}$. By Lemmas 2.2 and 3.2, α_{new1} has $o(n)$ bits and can be computed from S in $O(n)$ time.

- The height of the subtree rooted at v_i is $L[\text{index}_{\max}(L, \ell_i, r_i)]$ minus the depth of v_i in T .
- The lowest common ancestor v_k of v_i and v_j with $\ell_i < \ell_j$ can be determined as follows. If $r_i > r_j$, then $v_k = v_i$. Otherwise, $S[\text{index}_{\min}(L, r_i, \ell_j)]$ has to be a closed parenthesis r_x such that v_x is a child of v_k , as observed by Bender and Farach-Colton [2000].
- The distance of v_i and v_j is exactly the depth of v_i plus the depth of v_j minus two times of the depth of v_k , where v_k is the lowest common ancestor of v_i and v_j .

By Lemmas 2.2 and 3.2, the above queries can all be answered from S and α_{new1} in $O(1)$ time. \square

4. Rank and Select for Children

Before solving rank and select for children, we introduce the following definition and its property. A non-root node v_i is k -far if $w(\ell_p, \ell_i) > k$ and $w(\ell_i, r_p) > k$, where v_p is the parent of v_i .

LEMMA 4.1. *If v_i and v_j are two k -far non-root nodes with $|w(\ell_i, \ell_j)| \leq k$, then v_i and v_j are siblings.*

PROOF. Without loss of generality, we assume $\ell_i < \ell_j$. Since v_i and v_j are k -far non-root nodes with $w(\ell_i, \ell_j) \leq k$, v_i cannot be an ancestor or descendant of v_j . Thus, we have $r_i < \ell_j$. Assume for a contradiction that v_p (respectively, v_q) is the parent of v_i (respectively, v_j) and $v_p \neq v_q$. Observe that either $r_i < \ell_q$ or $r_p < \ell_j$ holds. Since v_j is k -far, $r_i < \ell_q$ implies $w(r_i, \ell_j) > k$. Since v_i is k -far,

$r_p < \ell_j$ implies $w(r_i, \ell_j) > k$. Either case leads to a contradiction, so the lemma is proved. \square

For presentational brevity, we classify non-root nodes into the following three disjoint classes: A node is

- narrow* if it is not *b*-far;
- medium* if it is *b*-far but not *B*-far; and
- wide* if it is *B*-far.

4.1. CHILD RANK. Let $child_rank(S, v_k)$ denote the number c such that v_k is the c th child of its parent. We have the following theorem.

THEOREM 4.2. *It takes $O(n)$ time to compute an $o(n)$ -bit string α_{new2} such that $child_rank(S, v_k)$ for each node v_k can be answered from S and α_{new2} in $O(1)$ time.*

PROOF. Let v_p be the parent of v_k . If $S[i, j]$ is a balanced string of parentheses, let $sibling(S, i, j)$ be the number of non-enclosed parenthesis pairs in $S[i, j]$. Observe that

$$\begin{aligned} child_rank(S, v_k) &= sibling(S, \ell_p + 1, \ell_k - 1) + 1 \\ &= degree(S, v_p) - sibling(S, \ell_k, r_p - 1) + 1. \end{aligned}$$

Therefore, it remains to support each query $sibling(S, i, j)$ in $O(1)$ time.

If v_k is narrow, we only need to answer $sibling(S, i, j)$ with $w(i, j) \leq b$. We simply build an $O(n)$ -time obtainable table M_1 to store the answers for any possible inputs. That is, let $M_1[S[i, i + b - 1]][j - i + 1] = sibling(S, i, j)$ for any indices i and j with $w(i, j) \leq b$. Since $sibling(S, i, j) \leq w(i, j)$, each entry requires $O(\log b)$ bits and M_1 takes $O(2^b b \log b) = o(n)$ bits.

If v_k is medium, we cannot afford to store all the answers of $sibling(S, i, j)$ with $w(i, j) \leq B$. We split S into length- b blocks. By Lemma 4.1, any two medium nodes v_i and v_j with $|w(\ell_i, \ell_j)| \leq b$ have the same parent, so for each block we save at most one medium node as a shortcut. Define tables M_2 and M_3 as follows. For each $t = 1, 2, \dots, n_b$,

- let $M_2[t] = (\ell_i, sibling(S, \ell_p + 1, \ell_i - 1))$, where ℓ_i is the smallest index, if any, with $(t - 1)b < \ell_i \leq tb$ such that v_i is a medium child of v_p with $w(\ell_p, \ell_i) \leq B$; and
- let $M_3[t] = (\ell_i, sibling(S, \ell_i, r_p - 1))$, where ℓ_i is the smallest index, if any, with $(t - 1)b < \ell_i \leq tb$ such that v_i is a medium child of v_p with $w(\ell_i, r_p) \leq B$.

Note that M_2 and M_3 have n_b entries, each requiring $O(\log B)$ bits, so both of them take $O(n_b \log B) = o(n)$ bits. Therefore, for any medium child v_k of v_p , if $w(\ell_p, \ell_k) \leq B$, then

$$\begin{aligned} sibling(S, \ell_p + 1, \ell_k - 1) &= sibling(S, \ell_p + 1, \ell_i - 1) + sibling(S, \ell_i, \ell_k - 1) \\ &= m + M_1[S[\ell_i, \ell_i + b - 1]][\ell_k - \ell_i], \end{aligned}$$

where $(\ell_i, m) = M_2[\lceil \frac{\ell_k}{b} \rceil]$. Similarly, if $w(\ell_k, r_p) \leq B$, then

$$\begin{aligned} sibling(S, \ell_k, r_p - 1) &= sibling(S, \ell_i, r_p - 1) - sibling(S, \ell_i, \ell_k - 1) \\ &= m - M_1[S[\ell_i, \ell_i + b - 1]][\ell_k - \ell_i], \end{aligned}$$

where $(\ell_i, m) = M_3[\lceil \frac{\ell_k}{b} \rceil]$.

function *child_rank*(S, v_k)

- 1: let v_p be the parent of v_k ;
- 2: **if** $w(\ell_p, \ell_k) \leq b$, **then** return $M_1[S[\ell_p + 1, \ell_p + b]][\ell_k - \ell_p - 1] + 1$;
- 3: **if** $w(\ell_k, r_p) \leq b$, **then** return $\text{degree}(S, v_p) - M_1[S[\ell_k, \ell_k + b - 1]][r_p - \ell_k] + 1$;
- 4: **if** $w(\ell_p, \ell_k) \leq B$, **then** let $(\ell_i, m) = M_2[\lceil \frac{\ell_k}{b} \rceil]$, and return $m + M_1[S[\ell_i, \ell_i + b - 1]][\ell_k - \ell_i] + 1$;
- 5: **if** $w(\ell_k, r_p) \leq B$, **then** let $(\ell_i, m) = M_3[\lceil \frac{\ell_k}{b} \rceil]$, and return $\text{degree}(S, v_p) - m + M_1[S[\ell_i, \ell_i + b - 1]][\ell_k - \ell_i] + 1$;
- 6: let $(\ell_j, m) = M_5[\lceil \frac{\ell_k}{B} \rceil][\lceil \frac{\ell_k \bmod B}{b} \rceil]$, and return $M_4[\lceil \frac{\ell_k}{B} \rceil] + m + M_1[S[\ell_j, \ell_j + b - 1]][\ell_k - \ell_j] + 1$;

FIG. 2. An $O(1)$ -time algorithm that computes *child_rank*(S, v_k).

Similar tricks work for wide nodes, but they have to be applied in two levels. We first split S into length- B blocks. For each $t = 1, 2, \dots, n_B$, let $M_4[t] = \text{sibling}(S, \ell_p + 1, \ell_i - 1)$, where ℓ_i is the smallest index, if any, with $(t - 1)B < \ell_i \leq tB$ such that v_i is a wide child of v_p . We further split each length- B block into length- b blocks. For each $t = 1, 2, \dots, n_B$ and $u = 1, 2, \dots, \lceil \frac{B}{b} \rceil$, let $M_5[t][u] = (\ell_j, \text{sibling}(S, \ell_p + 1, \ell_j - 1) - M_4[t])$, where ℓ_j is the smallest index, if any, with $(u - 1)b < \ell_j - (t - 1)B \leq ub$ such that v_j is a wide child of v_p . Note that $\text{sibling}(S, \ell_p + 1, \ell_j - 1) - M_4[t] \leq B$. One can easily verify that the number of bits required by M_4 is $O(n_B \log n) = o(n)$ and the number of bits required by M_5 is $O(n_B \lceil \frac{B}{b} \rceil \log B) = o(n)$. Thus, for any wide child v_k of v_p , we have

$$\begin{aligned} \text{sibling}(S, \ell_p + 1, \ell_k - 1) &= \text{sibling}(S, \ell_p + 1, \ell_j - 1) + \text{sibling}(S, \ell_j, \ell_k - 1) \\ &= M_4\left[\left\lceil \frac{\ell_k}{B} \right\rceil\right] + m + M_1[S[\ell_j, \ell_j + b - 1]][\ell_k - \ell_j], \end{aligned}$$

where $(\ell_j, m) = M_5[\lceil \frac{\ell_k}{B} \rceil][\lceil \frac{\ell_k \bmod B}{b} \rceil]$.

Finally, let $\alpha_{\text{new2}} = \alpha_{\text{aux}} \circ M_1 \circ M_2 \circ M_3 \circ M_4 \circ M_5$, which is an $o(n)$ -bit string obtainable from S in $O(n)$ time. The $O(1)$ -time algorithm for computing *child_rank*(S, v_k) is shown in Figure 2. \square

4.2. CHILD SELECT. First, we need the following lemmas to handle the select query for children. For any node v_i , let $\text{index}_c(S, \ell_i, m, c) = \ell_j - \ell_i$, where v_j is a sibling of v_i with $w(\ell_i, \ell_j) \leq m$ such that $\text{child_rank}(S, v_j) = \text{child_rank}(S, v_i) + c$. If such a v_j does not exist, $\text{index}_c(S, \ell_i, m, c) = \phi$.

LEMMA 4.3. *It takes $O(n)$ time to compute an $o(n)$ -bit string α_b such that $\text{index}_c(S, \ell_i, b^2, c)$ for any node v_i and index c can be computed from S and α_b in $O(1)$ time.*

PROOF. We simply build an $O(n)$ -time obtainable table α_b to store the answers for any possible inputs. That is, let $\alpha_b[S[\ell_i, \ell_i + b^2 - 1]][c] = \text{index}_c(S, \ell_i, b^2, c)$ for any node v_i and index c . Since each entry takes $O(\log b)$ bits, α_b requires $O(2^{b^2} b^2 \log b) = o(n)$ bits. \square

LEMMA 4.4. *Given a node v_i , it takes $O(B)$ time to compute an $o(B)$ -bit string $\alpha_B(\ell_i)$ such that $\text{index}_c(S, \ell_i, B, c)$ for any index c can be computed from S , α_b , and $\alpha_B(\ell_i)$ in $O(1)$ time.*

PROOF. For each $t = 0, 1, \dots, \lceil \frac{B}{b} \rceil - 1$, let $W_1[t] = \text{index}_c(S, \ell_i, B, tb)$. W_1 takes $O(\lceil \frac{B}{b} \rceil \log B) = o(B)$ bits. If $w(W_1[t], W_1[t + 1]) > b^2$, we save the answers

of $\text{index}_c(S, \ell_i, B, tb + z)$ for each $z = 0, 1, \dots, b - 1$ in W_2 . W_2 takes at most $O(\lceil \frac{B}{b^2} \rceil b \log B) = o(B)$ bits. Otherwise, by Lemma 4.3 $\text{index}_c(S, \ell_i, B, tb + z)$ can be computed in $O(1)$ time using $W_1[t] + \text{index}_c(S, \ell_i + W_1[t], b^2, z)$. Let $\alpha_B(\ell_i) = W_1 \circ W_2$, which has $o(B)$ bits and is obtainable in $O(B)$ time. \square

Given an array A of $\lceil \frac{m}{u} \rceil$ positive $\lceil \log u \rceil$ -bit integers with $m \leq n$ and $u = \lceil \log^3 m \rceil$, let $\text{index}_{\text{sum}}(A, x)$ denote the largest index y with $\sum_{t=1}^y A[t] < x$.

LEMMA 4.5. *It takes $O(m)$ time to compute an $o(m)$ -bit string $\alpha_A(A, m)$ such that $\text{index}_{\text{sum}}(A, x)$ for any index x can be determined from A and $\alpha_A(A, m)$ in $O(1)$ time.*

PROOF. This is a special case of the search query of the searchable partial sums problem [Raman et al. 2001; Hon et al. 2003]. Theorem 3 of Hon et al. [2003] gave an $o(m)$ -bit auxiliary string to support this query in $O(1)$ time, but it is unclear whether the preprocessing time is $O(m)$. Let us briefly prove this lemma as follows:

Let $d(x_1, x_2)$ denote $\text{index}_{\text{sum}}(A, x_2) - \text{index}_{\text{sum}}(A, x_1)$. For each $t = 0, \dots, \lceil \frac{m}{u} \rceil - 1$, let $W_3[t] = \text{index}_{\text{sum}}(A, tu)$. W_3 needs $O(\lceil \frac{m}{u} \rceil \log m) = o(m)$ bits. If $d(tu, (t+1)u) > \lceil \log^2 u \rceil$, for each $z = 0, 1, \dots, u - 1$, we save the values of $d(tu, tu + z)$ in W_4 . Because A is an array of positive integers, we have $d(tu, tu + z) \leq z$ and W_4 needs at most $O(\lceil \frac{m}{u \log^2 u} \rceil u \log u) = o(m)$ bits. Otherwise, let

$$W_5[A[\text{index}_{\text{sum}}(A, tu), \text{index}_{\text{sum}}(A, tu) + \lceil \log^2 u \rceil - 1]][z] = d(tu, tu + z)$$

for each $z = 0, 1, \dots, u - 1$. W_5 takes $O(2^{\log^3 u} u \log \log u) = o(m)$ bits and is obtainable in $O(m)$ time. Now, let $\alpha_A(A, m) = W_3 \circ W_4 \circ W_5$, which requires $o(m)$ bits and can be obtained in $O(m)$ time. To answer $\text{index}_{\text{sum}}(A, x)$ in $O(1)$ time, first let t and z be the integers with $x = tu + z$ and $0 \leq z < u$, and then find the values of $\text{index}_{\text{sum}}(A, tu)$ and $d(tu, tu + z)$ from $\alpha_A(A, m)$. The answer is $\text{index}_{\text{sum}}(A, tu) + d(tu, tu + z)$. \square

Let $\text{child_select}(S, v_p, c)$ denote the index ℓ_k such that v_k is the c th child of v_p . We have the following theorem.

THEOREM 4.6. *It takes $O(n)$ time to compute an $o(n)$ -bit string α_{new3} such that $\text{child_select}(S, v_p, c)$ for each node v_p and c can be answered from S and α_{new3} in $O(1)$ time.*

PROOF. We say that nodes in a set D are d -disjoint [Chiang et al. 2005] if

- $w(\ell_i, r_i) > d$ holds for any node v_i in D ; and
- any two nodes v_i and v_j in D satisfy at least one of $|w(\ell_i, \ell_j)| > d$ and $|w(r_i, r_j)| > d$.

Let X be a $2\lceil \frac{2n}{d} \rceil$ -element array. For each $t = 1, 2, \dots, \lceil \frac{2n}{d} \rceil$, we store v_i in $X[2t - 1]$, where ℓ_i is the smallest index, if any, with $(t-1)d < \ell_i \leq td$ such that v_i is in D ; and also store v_j in $X[2t]$, where r_j is the largest index, if any, with $(t-1)d < r_j \leq td$ such that v_j is in D . Then, every node v_i in D takes at least one slot in X , and can be easily verified using ℓ_i and r_i . We simply say that X has v_i if and only if v_i takes at least one of $X[2\lceil \frac{\ell_i}{d} \rceil - 1]$ or $X[2\lceil \frac{r_i}{d} \rceil]$. For notational brevity, let $X[v_i]$ denote the element taken by v_i .

The preprocessing is under the following traversal procedure: first traverse each node v_p of T in prefix order, and for each v_p traverse every child v_i of v_p in counterclockwise order. Since selecting and matching a parenthesis on S takes $O(1)$ time, and each node is traversed at most two times, one as v_p and the other as v_i , the whole procedure takes $O(n)$ time. The discussion below focuses on nodes v_p and v_i in each iteration of the aforementioned traversal.

—*Case 1.* v_i is a wide child of v_p . Let *counter* denote the number of wide nodes discovered before each iteration. It is not difficult to see that the parents of wide nodes are B -disjoint. Let X_1 be the $2n_B$ -element array with $X_1[v_p] = (\text{before}_p, \text{first}, \text{last})$, where before_p is the value of *counter* before we get v_p , and *first* (respectively, *last*) is the rank of the first (respectively, last) wide child of v_p . Then we partition S into length- B blocks. Let Y_1 be the n_B -element array with $Y_1[t] = (\text{before}_i, \ell_i)$, where ℓ_i is the smallest index in a block such that v_i is wide, before_i is the value of *counter* before we get v_i , and t is the first empty entry of Y_1 . Both of X_1 and Y_1 take $O(n_B \log n) = o(n)$ bits.

—*Case 2.* v_i is a medium child of v_p . First, we partition S into length- B blocks.

If $w(\ell_p, \ell_i) \leq B$, we say that v_i belongs to the $\lceil \frac{\ell_p}{B} \rceil$ -th block, otherwise the $\lceil \frac{r_p}{B} \rceil$ -th block. For each $t = 1, 2, \dots, n_B$, let *counter*[t] denote the number of medium nodes belonging to the t th block before each iteration. Note that at most B medium nodes belong to a block. Similarly, one can verify that the parents of medium nodes are b -disjoint. Let X_2 be the $2n_b$ -element array with $X_2[v_p] = (\text{before}_L, \text{first}_L, \text{last}_L, \text{before}_R, \text{first}_R, \text{last}_R)$, where

—*before* _{L} (respectively, *before* _{R}) is the value of *counter*[$\lceil \frac{\ell_p}{B} \rceil$] (respectively, the value of *counter*[$\lceil \frac{r_p}{B} \rceil$]) before we get v_p ,

—*first* _{L} (respectively, *first* _{R}) is the rank of the first medium child of v_p belonging to the $\lceil \frac{\ell_p}{B} \rceil$ -th (respectively, $\lceil \frac{r_p}{B} \rceil$ -th) block, and

—*last* _{L} (respectively, *last* _{R}) is the rank of the last medium child of v_p belonging to the $\lceil \frac{\ell_p}{B} \rceil$ -th (respectively, $\lceil \frac{r_p}{B} \rceil$ -th) block.

Note that $1 \leq \text{first}_L \leq \text{last}_L \leq B$ and $\text{degree}(S, v_p) - B \leq \text{first}_R \leq \text{last}_R \leq \text{degree}(S, v_p)$. We further partition each length- B block into length- b blocks.

For each $t = 1, 2, \dots, n_B$, let $Y_2[t]$ be the $\lceil \frac{B}{b} \rceil$ -element array with $Y_2[t][u] = (\text{before}_i, \ell_i)$, where ℓ_i is the smallest index in a length- b block such that v_i is a medium node belonging to the t th length- B block, *before* is the value of *counter*[t] before we get v_k , and u is the first empty entry of $Y_2[t]$. Observe that X_2 needs $O(n_b \log B) = o(n)$ bits and Y_2 needs $O(n_B \lceil \frac{B}{b} \rceil \log B) = o(n)$ bits.

For each $t = 1, 2, \dots, n_B$, let $\alpha_{B1}[t] = \alpha_B(\ell_i)$ with $(\text{before}_i, \ell_i) = Y_1[t]$. By Lemma 4.4, α_{B1} takes $o(n)$ bits and is obtainable in $O(n)$ time. Let A_1 be the n_B -element array such that $\sum_{i=1}^u A_1[t] = \text{before}_i$ with $(\text{before}_i, \ell_i) = Y_1[u]$ holds for each $u = 1, 2, \dots, n_B$. Note that $0 < A_1[t] \leq B$ holds for any index t , so A_1 takes $O(n_B \log B) = o(n)$ bits. Also, for each $t = 1, 2, \dots, n_B$, let $A_2[t]$ be the $\lceil \frac{B}{b} \rceil$ -element array such that $\sum_{u=1}^x A_2[t][u] = \text{before}_i$ with $(\text{before}_i, \ell_i) = Y_2[t][x]$ holds for each $x = 1, 2, \dots, \lceil \frac{B}{b} \rceil$. Observe that $0 < A_2[t][u] \leq b$ holds for any indices t and u , so A_2 takes $O(n_B \lceil \frac{B}{b} \rceil \log b) = o(n)$ bits. Let $\alpha_{A1} = \alpha_A(A_1, n)$, and for each $t = 1, 2, \dots, n_B$, let $\alpha_{A2}[t] = \alpha_A(A_2[t], B)$. By Lemma 4.5, both of α_{A1} and α_{A2} take $o(n)$ bits and are obtainable in $O(n)$ time. At last, we construct an

```

function child_select( $S, v_p, c$ )
1: if  $X_1$  has  $v_p$  then
2:   let  $(before_p, first, last) = X_1[v_p]$ ;
3:   if  $first \leq c \leq last$  then
4:     let  $z = before_p + c - first + 1$  and  $(before_i, \ell_i) = Y_1[index_{sum}(A_1, z)]$ ;
5:     return  $\ell_i + index_c(S, \ell_i, B, z - before_i)$ ;
6:   end if
7: end if
8: if  $X_2$  has  $v_p$  then
9:   let  $(before_L, first_L, last_L, before_R, first_R, last_R) = X_2[v_p]$ ;
10:  if  $first_L \leq c \leq last_L$  then
11:    let  $t = \lceil \frac{\ell_p}{B} \rceil$ ,  $z = before_L + c - first_L + 1$ , and  $(before_i, \ell_i) = Y_2[t][index_{sum}(A_2[t], z)]$ ;
12:    return  $\ell_i + index_c(S, \ell_i, b^2, z - before_i)$ ;
13:  end if
14:  if  $first_R \leq c \leq last_R$  then
15:    let  $t = \lceil \frac{r_p}{B} \rceil$ ,  $z = before_R + c - first_R + 1$ , and  $(before_i, \ell_i) = Y_2[t][index_{sum}(A_2[t], z)]$ ;
16:    return  $\ell_i + index_c(S, \ell_i, b^2, z - before_i)$ ;
17:  end if
18: end if
19: if  $index_c(S, \ell_p + 1, b^2, c) \neq \phi$ , then return  $\ell_p + 1 + index_c(S, \ell_p + 1, b^2, c)$ ;
20: else return  $r_p - F[S[r_p - b + 1, r_p]][degree(S, v_p) - c]$ ;

```

FIG. 3. An $O(1)$ -time algorithm that computes $child_select(S, v_p, c)$.

$O(n)$ -time obtainable table F with $F[S[r_p - b + 1, r_p]][degree(S, v_p) - c] = r_p - \ell_i$, where v_i is the c th child of v_p with $w(\ell_i, r_p) \leq b$. Note that $degree(S, v_p) - c \leq b$, so F takes $O(2^b b \log b) = o(n)$ bits.

To implement $child_select$ in $O(1)$ time, let $\ell_k = child_select(S, v_p, c)$. v_k is wide if and only if X_1 has v_p and $first \leq c \leq last$, where $(before_p, first, last) = X_1[v_p]$. Moreover, letting $z = before_p + c - first + 1$, v_k is the z th wide node discovered during the traversal procedure. Let $(before_i, \ell_i) = Y_1[index_{sum}(A_1, z)]$, so v_k is a sibling of v_i with $w(\ell_i, \ell_k) \leq B$ such that $child_rank(S, v_k) = child_rank(S, v_i) + z - before_i$. By Lemma 4.4, we can locate v_k using $\ell_k = \ell_i + index_c(S, \ell_i, B, z - before_i)$.

v_k is medium if and only if X_2 has v_p and at least one of $first_L \leq c \leq last_L$ and $first_R \leq c \leq last_R$ is satisfied, where $(before_L, first_L, last_L, before_R, first_R, last_R) = X_2[v_p]$. If $first_L \leq c \leq last_L$, let $t = \lceil \frac{\ell_p}{B} \rceil$ and $z = before_L + c - first_L + 1$. If $first_R \leq c \leq last_R$, let $t = \lceil \frac{r_p}{B} \rceil$ and $z = before_R + c - first_R + 1$. Then, v_k is the z th medium node belonging to the t th length- B block discovered during the traversal procedure. Let $(before_i, \ell_i) = Y_2[t][index_{sum}(A_2[t], z)]$, so v_k is a sibling of v_i with $w(\ell_i, \ell_k) \leq b$ such that $child_rank(S, v_k) = child_rank(S, v_i) + z - before_i$. By Lemma 4.3, we can locate v_k using $\ell_k = \ell_i + index_c(S, \ell_i, b^2, z - before_i)$.

If v_k is neither wide nor medium, it must be narrow. If $index_c(S, \ell_p + 1, b^2, c) \neq \phi$, then we have $\ell_k = \ell_p + 1 + index_c(S, \ell_p + 1, b^2, c)$. Otherwise, $\ell_k = r_p - F[S[r_p - b + 1, r_p]][degree(S, v_p) - c]$.

Finally, let $\alpha_{new3} = \alpha_{aux} \circ \alpha_b \circ \alpha_{B1} \circ X_1 \circ Y_1 \circ X_2 \circ Y_2 \circ A_1 \circ \alpha_{A1} \circ A_2 \circ \alpha_{A2} \circ F$, which takes $o(n)$ bits and can be computed from S in $O(n)$ time. The $O(1)$ -time algorithm for computing $child_select(S, v_p, c)$ is shown in Figure 3. \square

ACKNOWLEDGMENT. We thank Kai-min Chung for helpful discussion. We also thank the anonymous reviewers for their helpful comments.

REFERENCES

- BELL, T. C., CLEARY, J. G., AND WITTEN, I. H. 1990. *Text Compression*. Prentice-Hall, Englewood Cliffs, NJ.
- BENDER, M. A., AND FARACH-COLTON, M. 2000. The LCA problem revisited. In *Proceedings of the 4th Latin American Symposium on Theoretical Informatics* (Punta del Este, Uruguay), G. H. Gonnet, D. Panario, and A. Viola, Eds. Lecture Notes in Computer Science, vol. 1776. Springer-Verlag, New York, 88–94.
- BENOIT, D., DEMAINE, E. D., MUNRO, J. I., RAMAN, R., RAMAN, V., AND RAO, S. S. 2005. Representing trees of higher degree. *Algorithmica* 43, 4, 275–292.
- BONICHON, N., GAVOILLE, C., HANUSSE, N., POULALHON, D., AND SCHAEFFER, G. 2006. Planar graphs, via well-orderly maps and trees. *Graph and Combinat.* 22, 1–18.
- CHIANG, Y.-T., LIN, C.-C., AND LU, H.-I. 2005. Orderly spanning trees with applications. *SIAM J. Comput.* 34, 4, 924–945.
- CHUANG, R. C.-N., GARG, A., HE, X., KAO, M.-Y., AND LU, H.-I. 1998. Compact encodings of planar graphs via canonical ordering and multiple parentheses. In *Proceedings of the 25th International Colloquium on Automata, Languages, and Programming* (Aalborg, Denmark), K. G. Larsen, S. Skyum, and G. Winskel, Eds. Lecture Notes in Computer Science, vol. 1443. Springer-Verlag, New York, 118–129.
- CLARK, D. R. 1996. Compact PAT trees. Ph.D. dissertation, University of Waterloo.
- CLARK, D. R., AND MUNRO, J. I. 1996. Efficient suffix trees on secondary storage. In *Proceedings of the 7th Annual ACM-SIAM Symposium on Discrete Algorithms* (Atlanta, CA). ACM, New York, 383–391.
- DELPRAAT, O., RAHMAN, N., AND RAMAN, R. 2006. Engineering the LOUDS succinct tree representation. In *Proceedings of the 5th International Workshop on Experimental Algorithms* (Cala Galdana, Menorca, Spain), Lecture Notes in Computer Science, vol. 4007. Springer-Verlag, New York, 134–145.
- ELIAS, P. 1975. Universal codeword sets and representations of the integers. *IEEE Trans. Inf. Theory IT-21*, 194–203.
- GABOW, H. N., BENTLEY, J. L., AND TARJAN, R. E. 1984. Scaling and related techniques for geometry problems. In *Proceedings of the 16th Annual ACM Symposium on Theory of Computing*. ACM, New York, 135–143.
- GEARY, R. F., RAHMAN, N., RAMAN, R., AND RAMAN, V. 2004a. A simple optimal representation for balanced parentheses. In *Proceedings of the 15th Annual Symposium on Combinatorial Pattern Matching* (Istanbul, Turkey), S. C. Sahinalp, S. Muthukrishnan, and U. Dogrusöz, Eds. Lecture Notes in Computer Science, vol. 3109. Springer-Verlag, New York, 159–172.
- GEARY, R. F., RAMAN, R., AND RAMAN, V. 2004b. Succinct ordinal trees with level-ancestor queries. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms* (New Orleans, LA), J. I. Munro, Ed. ACM, New York, 1–10.
- GRAHAM, R. L., KNUTH, D. E., AND PATASHNIK, O. 1989. *Concrete Mathematics*. Addison-Wesley, Reading, MA.
- HE, X., KAO, M.-Y., AND LU, H.-I. 1999. Linear-time succinct encodings of planar graphs via canonical orderings. *SIAM J. Disc. Math.* 12, 3, 317–325.
- HON, W.-K., SADAKANE, K., AND SUNG, W.-K. 2003. Succinct data structures for searchable partial sums. In *Proceedings of the 14th Symposium on Algorithms and Computation* (Kyoto, Japan), T. Ibaraki, N. Katoh, and H. Ono, Eds. Lecture Notes in Computer Science, vol. 2906. Springer-Verlag, New York, 505–516.
- JACOBSON, G. 1989. Space-efficient static trees and graphs. In *Proceedings of the 30th Annual Symposium on Foundations of Computer Science* (Research Triangle Park, NC). IEEE Computer Society Press, Los Alamitos, CA, 549–554.
- MUNRO, J. I. 1996. Tables. In *Proceedings of the 16th Conference on Foundations of Software Technology and Theoretical Computer Science* (Hyderabad, India), Lecture Notes in Computer Science, vol. 1180. Springer-Verlag, New York, 37–42.
- MUNRO, J. I., AND RAMAN, V. 2001. Succinct representation of balanced parentheses, static trees and planar graphs. *SIAM J. Comput.* 31, 3, 762–776.
- MUNRO, J. I., RAMAN, V., AND RAO, S. S. 2001. Space efficient suffix trees. *J. Algor.* 39, 2, 205–222.
- MUNRO, J. I., AND RAO, S. S. 2004. Succinct representations of functions. In *Proceedings of the 31st International Colloquium on Automata, Languages and Programming* (Turku, Finland), Lecture Notes in Computer Science, vol. 3142. Springer-Verlag, New York, 1006–1015.

- RAMAN, R., RAMAN, V., AND RAO, S. S. 2001. Succinct dynamic data structures. In *Proceedings of the 7th International Workshop on Algorithms and Data Structures* (Providence, RI), F. K. H. A. Dehne, J.-R. Sack, and R. Tamassia, Eds. Lecture Notes in Computer Science, vol. 2125. Springer-Verlag, New York, 426–437.
- SADAKANE, K. 2002. Succinct representations of LCP information and improvements in the compressed suffix arrays. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms*, (San Francisco, CA). ACM, New York, 225–232.
- VAN EMDE BOAS, P. 1990. Machine models and simulations. In *Handbook of Theoretical Computer Science*, J. van Leeuwen, Ed. Vol. A. Elsevier, Amsterdam, The Netherlands, Chapter 1, 1–60.

RECEIVED DECEMBER 2005; REVISED OCTOBER 2006 AND MARCH 2007; ACCEPTED APRIL 2007