Correctness of Constructing Optimal Alphabetic Trees Revisited

Marek Karpinski *
Lawrence L. Larmore †
Wojciech Rytter ‡

TR-96-053

December 1996

Abstract

Several new observations which lead to new correctness proofs of two known algorithms (Hu-Tucker and Garsia-Wachs) for construction of optimal alphabetic trees are presented. A generalized version of the Garsia-Wachs algorithm is given. Proof of this generalized version works in a structured and illustrative way and clarifies the usually poorly-understood behavior of both the Hu-Tucker and Garsia-Wachs algorithms. The generalized version permits any non-negative weights, as opposed to strictly positive weights required in the original Garsia-Wachs algorithm. New local structural properties of optimal alphabetic trees are given. The concept of *well-shaped*

*Dept. of Computer Science, University of Bonn, 53117 Bonn, and the International Computer Science Institute, Berkeley. Research partially done while visiting Dept. of Computer Science, Princeton University. Research supported by DFG Grant KA 673/4-1, and by the ESPRIT BR Grants 7097 and EC-US 030 and by DIMACS. Email: marek@cs.uni-bonn.de

†Department of Computer Science, University of Nevada, Las Vegas, NV 89154–4019, USA. Partially supported by National Science Foundation grants CCR-9112067 and CCR-9503441. Email:larmore@cs.unlv.edu

‡Dept. of Computer Science, University of Warsaw, 02–097 Warszawa. Supported by the grant KBN 8T11C01208. Email:rytter@mimuw.edu.pl

segment (a part of an optimal tree) is introduced. It is shown that some parts of the optimal tree are known in advance to be well-shaped, and this implies correctness of the algorithms rather easily. The crucial part of the correctness proof of the Garsia-Wachs algorithm, namely the *structural theorem*, is identified. The correctness proof of the Hu-Tucker algorithm consists of showing a very simple mutual simulation between this algorithm and the Garsia-Wachs algorithm. For this proof, it is essential to use the generalized version of Garsia-Wachs algorithm, in which an arbitrary locally minimal pair is processed, not necessarily the rightmost minimal pair. Such a generalized version is also needed for parallel implementations. Another result presented in this paper is the clarification of the problem of resolving ties (equalities between weights of items) in the Hu-Tucker algorithm. This is related to the proof, by simulation, of correctness of the Hu-Tucker algorithm. It is shown that the condition that there are no ties may generally be assumed without harm and that, essentially, the Hu-Tucker algorithm avoids ties automatically.

1 Introduction

Recently, there has been a renewed interest in the problem of construction of optimal alphabetic trees [6, 10, 11, 12]. The Hu-Tucker (HT) algorithm [3] is a celebrated classical algorithm for this problem, whose correctness is not widely understood. The Garsia-Wachs (GW) algorithm [1], has a simpler but still very technical proof based on several formal claims proved simultaneously by induction. Our proof of correctness of the HT algorithm works by reducing to correctness of a general version of the GW algorithm in which *any* locally minimal pair is processed, not necessarily the rightmost one. This general version is also needed in our parallel implementations (see [11]). A restricted version of the GW algorithm (for *rightmost* minimal pairs) was considered in [5].

The aim of this paper is to provide proofs of correctness of both the HT and the GW algorithms that are more structural than those in the original papers. The simplicity of the description of both algorithms is misleading. The original correctness proofs are very intricate. According to Knuth, "No simple proof is known, and it is quite possible that no simple proof will ever be found!" ([7], page 443.) We provide several new facts about the local structure of optimal alphabetic trees, introducing new local operations on trees, and specify a mutual simulation between both algorithms.

Statement of the optimal alphabetic tree problem. Assume we have n weighted items, where p_i is the non-negative *weight* of the i^{th} item. Write $\alpha = p_1 \dots p_n$.

The Garsia-Wachs algorithm permutes α. We adopt the convention that the items of α have unique names, and that these names are preserved when items are moved. When convenient to do so, we will assume that those names are the positions of items in the list, namely integers in $[1 \dots n]$.
An alphabetic tree over α is an ordered binary tree T with n leaves, where the i^{th} leaf (in left-to-right order) corresponds to the i^{th} item of α. Throughout this paper, a binary tree must be *full*, *i.e.*, each internal node must have exactly two sons. We define the *cost* of any alphabetic tree T as follows:

$$cost(T) = \sum_{i=1}^{n} p_i level_T(i)$$

where $level_T$ is the *level function* of T, *i.e.*, $level_T(i)$ is the level (or depth) of i in T, defined to be the length of the path in T from the root to i. The optimal alphabetic tree problem (OAT problem) is to find an alphabetic tree of minimum cost. Both the GW and HT algorithms have two phases. The first phase constructs the level function $level_T$ of an optimal alphabetic tree T. The second phase constructs T from its level function. a relatively trivial procedure that takes linear time. In fact, throughout the rest of the paper,

we ignore this second phase, and take the array of values of the level function to be the output of any algorithm for the OAT problem.

Construction of the optimal alphabetic tree. The alphabetic tree is constructed by reducing the initial sequence of items to a shorter sequence in a manner similar to that of the Huffman algorithm, with one important difference. In the Huffman algorithm, the minimum pair of items are combined, because it can be shown that they are siblings in the optimal tree. If we could identify two adjacent items that are siblings in the optimal alphabetic tree, we could combine them and then proceed recursively. Unfortunately, there is no known way to identify such a pair. Even a minimal pair may not be siblings. Consider the weight sequence (8 7 7 8). The second and the third items are not siblings in any optimal alphabetic tree.

Instead, the HT and GW algorithms, as well as the algorithms of [6, 10, 11, 12], operate by identifying a pair of items that have the same level in the optimal tree. These items are then combined into a single "package," reducing the number of items by one. The details on how this process proceeds differ in the different algorithms.

2 Correctness of the Garsia-Wachs algorithm

Define $TwoSum(i) = p_i + p_{i+1}$, the i^{th} *two-sum*, for $1 \leq i < n$. A pair of adjacent items $(i, i+1)$ is a *locally minimal pair* (or *lmp* for short) if

$$\begin{array}{rcll} TwoSum(i-1) & \geq & TwoSum(i) & \text{if} \quad i > 1 \\ TwoSum(i) & < & TwoSum(i+1) & \text{if} \quad i \leq n-2 \end{array}$$

A locally minimal pair which is currently being processed is called the *active pair*.

The Operator *Move*. If w is any item in a list π of weighted items, define $RightPos(w)$ to be the predecessor of the nearest right larger neighbor of w. In this context, "larger" means "greater than or equal to." If w has no right larger neighbor, define $RightPos(w)$ to be the last item of π. Let $Move(w, \pi)$ be the operator that changes π by moving w just to the right of $RightPos(w)$. Note that if $RightPos(w) = w$, then $Move(w, \pi)$ does nothing.

Similarly, if u, v are adjacent items in π, define $RightPos(u, v)$ to be the predecessor of the nearest item to the right of v whose weight is at least $weight(u) + weight(v)$. If there is no such item, define $RightPos(u, v)$ to be the last item of π. Let $Move(u, v, \pi)$ be the operator that changes π by moving u and v to just to the right of $RightPos(u, v)$. For

example, if $\pi = (1, \ldots n)$, and if $RightPos(i, i+1) = j$, then $Move(i, i+1, \pi)$ changes π to

$$\pi_{i,j} = (1, \ldots, i-1, i+2, \ldots, j, i, i+1, j+1, \ldots, n)$$

Two binary trees T_1 and T_2 are said to be **level equivalent** (we write $T_1 \cong T_2$) if T_1, and T_2 have the same set of leaves (possibly in a different order) and $level_{T_1} = level_{T_2}$.

Theorem 2.1 (Correctness of the GW algorithm) *Let $(i, i+1)$ be a locally minimal pair and $RightPos(i, i+1) = j$, and let T' be a tree over the sequence $\pi_{i,j}$, optimal among all trees over $\pi_{i,j}$ in which i, $i+1$ are siblings. Then there is an optimal alphabetic tree T over the original sequence $\pi = (1, \ldots n)$ such that $T \cong T'$.*

The significance of Theorem 2.1 is that $level_{T'}$ may be computed by combining i and $i+1$ into a single node, v, and then applying the procedure recursively on the resulting list of length $(n-1)$. Then $level_T(i) = level_T(i+1) = level_{T'}(v) + 1$, while $level_T = level_{T'}$ on all other items.

The array *level* is global of size $(2n-1)$. Its indices are the names of the nodes, *i.e.*, the original n items and the $(n-1)$ nodes ("packages") created during execution of the algorithm. The algorithm works in quadratic time, if implemented in a naive way. Using priority queues, it works in $O(n \log n)$ time. Correctness follows directly from Theorem 2.1.

```
procedure GW(π); {π is a sequence of names of items}
{General version of the Garsia-Wachs algorithm}
if π = (v) then
    level[v] = 0  else begin
      (*) find any locally minimal pair (u, w) of π
      create a new item v whose weight is p_u + p_w;
      replace u by the item v and delete w;
      (#) Move(v, π);
       GW(π);
      level[u] := level[w] := level[v] + 1;
    end;
```

Denote by $\mathrm{OPT}(i)$ the set of all alphabetic trees over the leaf-sequence $(1, \ldots n)$ which are optimal among trees in which i and $i+1$ are at the same level. Assume the pair $(i, i+1)$ is locally minimal. Let $\mathrm{OPT}_{moved}(i)$ be the set of all alphabetic trees over the leaf-sequence $\pi_{i,j}$ which are optimal among all trees in which leaves i and $i+1$ are at the same level, where $j = RightPos(i, i+1)$.

Two sets of trees OPT and OPT$'$ are said to be *level equivalent*, written OPT $\cong$ OPT$'$, if, for each tree $T \in$ OPT, there is a tree $T' \in$ OPT$'$ such that $T' \cong T$, and vice versa.

Theorem 2.2
Let $(i, i+1)$ be a locally minimal pair. Then

(1) $\mathrm{OPT}(i) \cong \mathrm{OPT}_{moved}(i)$.

(2) $\mathrm{OPT}(i)$ *contains an optimal alphabetic tree T.*

(3) $\mathrm{OPT}_{moved}(i)$ *contains a tree T' with i, $i+1$ as siblings.*

Theorem 2.2 directly implies Theorem 2.1. Points (2) are (3) are simple. We prove them in this section for completeness. Point (1) is rather subtle, for if we drop the requirement that $i, i+1$ are at the same level, then this point is false for some weight sequences, *e.g.*, (7 8 13 14 1).

Our main contribution is the discovery and a structural proof of Point (1).

Description of the shift operations. We introduce two useful local operations, *RightShift* and *LeftShift* on trees. Both operations change the shape of an alphabetic tree locally without changing the order of items (as leaves). We describe in detail only the operation *LeftShift*, as *RightShift* is similar. Assume $v_1, v_2, \ldots, v_k$ are roots of disjoint subtrees $T_1, T_2, \ldots, T_k$, for $k \geq 2$, and the segments of leaves covered by these subtrees are disjoint and cover (left to right) a segment of consecutive leaves. We call such a sequence of nodes $(v_1, \ldots, v_k)$ a *cut*.

Let $p(v)$ denote the father of any node v in T, and $p'(v)$ denote the father of v in T', the tree resulting from applying the operation *LeftShift* to the cut. Create a new node w. Let $p'(w) = p(v_1)$. Let $p'(v_1) = p'(v_2) = w$. For $2 < t < k$, let $p'(v_t) = p(v_{t-1})$. To finish the description, we consider two cases. If v_k is the right son of $p(v_k)$, let $p'(v_k) = p(p(v_k))$. Otherwise, let v_{k+1} be the sibling of v_k, let $p'(v_k) = p(v_{k-1})$ and let $p'(v_{k+1}) = p(p(v_k))$. In both cases, $p(v_k)$ is deleted from the tree. For all other nodes, $p'(v) = p(v)$.

Proof of point (2) of theorem 2.2. Assume the levels of i and $i+1$ are different in some optimal tree T, hence they are not siblings in T. If $level_T(i) < level_T(i+1)$ then we can perform *LeftShift*$(i, i+1)$, obtaining a new tree T'. The level of i increases by 1 and the level of $i+1$ does not increase. On the other hand the level of $i+2$ decreases at least by one. Hence (since $p_{i+2} \geq p_i$ and $p_{i+1} \geq 0$) $cost(T') \leq cost(T)$, i and $i+1$ are siblings in T' and their levels are equal, and thus $T' \in \mathrm{OPT}(i)$. If $level_T(i) > level_T(i+1)$ the proof is similar; use *RightShift*$(i, i+1)$.

Proof of point (3) of theorem 2.2. Consider a tree $T \in \mathrm{OPT}_{moved}(i)$. If $i, i+1$ are not siblings then $1 < i < n-1$. After applying *LeftShift*$(i, i+1)$ they become siblings. Both

of them go down but item $j+1$ goes one level up. Since $p_i + p_{i+1} - p_{j+1} \leq 0$ the resulting tree T' is still optimal.

3 The structural theorem

This section is devoted to the proof of Point (1) of theorem 2.2.

Informal overview of the proof of point (1) of correctness theorem. The crucial point is to show that the certain parts of trees in $\mathrm{OPT}(i)$ and $\mathrm{OPT}_{moved}(i)$ which are active with respect to the pair $(i, i+1)$ are "well-shaped" (in the sense defined below) and that this guarantees that the pair $(i, i+1)$ can be moved to the other side of such a part without affecting the level function. The point (1) of the correctness theorem is broken into the proof of the *movability lemma* and that of the *structural theorem*. The movability lemma is rather obvious. The structural theorem is proved by considering conditions of well-shaped segments and several cases. The proofs are by contradictions: if a certain condition is not satisfied for the optimal tree then using shift operations the tree is transformed into a tree of a smaller cost. This contradicts the optimality of the original tree. Point (1) of Theorem 2.2, and correctness of the GW algorithm, follow directly from Theorem 3.2 and Lemma 3.1.

Definition of well-shaped segment. Let $\mathrm{LCA}_T(u, w)$ denote the lowest common ancestor of nodes u and w in T.

We say that a set S of leaves of T is *h-isolated* if

1. For any $u \in S$, $level_T(u) \geq h$.
2. For any $u \in S$ and $w \notin S$, $level_T(\mathrm{LCA}(u, w)) \leq h$.

We say that a segment $[i \ldots j]$ of consecutive items (leaves) is *left well-shaped* at level h in T if $\{i, \ldots j\}$ is h-isolated and $level_T(i) = level_T(i+1) = h+1$. We define *right well-shaped* similarly (in this case $level_T(j-1) = level_T(j) = h+1$).

The leaves in the segment $[i \ldots j]$ and all their ancestors at level at least h is called the *active window*. Note that the active window is a forest.

The window is said to be well-shaped iff the sequence of its leaves is left or right well-shaped. The introduction of windows is useful in visualizing local properties and rearrangements, as these rearrangements occur inside such windows. Trees in $\mathrm{OPT}(i)$ and

$\mathrm{OPT}_{moved}(i)$ are illustrated in Figure 1 for $i = 3$ and for the weight sequence

(80 12 10 11 13 3 4 5 9 8 7 25)

The windows in trees T and T' are indicated by dotted lines. A window is also illustrated in Figure 2.

Lemma 3.1 (Movability lemma) *If the segment $[i \ldots j]$ is left well-shaped, then the active pair of items $(i, i+1)$ can be moved to the other side of the segment by locally rearranging subtrees in the active window without changing the relative order of the other items and without changing the level function of the tree.*

Proof. The proof is straightforward. Let $h = level(i) - 1$. There are four cases, depending on whether i and $i+1$ are siblings, and on whether $\mathrm{LCA}(j, j+1) = h$. Figure 2 illustrates the proof in one case. We omit the details. ∎

If v is a node in a given alphabetic tree then we write p_v for the total weight of the leaves of the subtree rooted at v.

Theorem 3.2 (Structural theorem)
Assume $(i, i+1)$ is an lmp, $j = RightPos(i, i+1)$, $T \in \mathrm{OPT}(i)$, and $T' \in \mathrm{OPT}_{moved}(i)$. Then (i) *the segment $[i+2, \ldots, j, i, i+1]$ is right well-shaped in T', and* (ii) *the segment $[i \ldots j]$ is left well-shaped in T.*

Proof. We shall initially assume that all weights are positive.

If $j = i+1$, the theorem is trivial. Thus, without loss of generality, $j > i+1$.

Let $h = level_T(i) - 1 = level_T(i+1) - 1$ and $h' = level_{T'}(i) - 1 = level_{T'}(i+1) - 1$.

Claim A: : For any $u \in [i+2 \ldots j]$ we have

(1) $level_T(u) \geq h$ and (2) $level_{T'}(u) \geq h'$.

Proof of Claim A. We show only the proof of point (1), as (2) has a very similar proof. The proof is by contradiction. Suppose the claim is false. Let k be the leftmost item such that $level_T(k) < h$. Let $T'' \in \mathrm{OPT}(i)$ be the tree obtained from T by applying $RightShift(i, i+1, v_1, \ldots, v_r, k)$, where $v_2, \ldots, v_r$ are at level h and v_1 is at level h if i and $i+1$ are siblings in T, level $h+1$ otherwise. If i and $i+1$ are siblings in T, i and $i+1$ go up and only k goes down (see Figure 3). Otherwise, $i-1$ and $i+2$ go up and only k goes down. Since $p_k < p_i + p_{i+1} < p_{i-1} + p_{i+2}$, T'' has lower cost than T, a contradiction.

Claim B: : If $j < n$, then $level_T(\mathrm{LCA}_T(j, j+1) \leq h$.
Proof of Claim B. The proof is by contradiction. Suppose the claim is false. Let w be the ancestor of j and $j+1$ at level $h+1$. Let w_1 and w_2 be the sons of w.

Case 1: There exists a leaf k at level $\leq h$ in the segment. Pick the rightmost such k. Let $T'' \in \mathrm{OPT}(i)$ be the tree obtained from T by applying *LeftShift*$(k, v_1, \ldots, v_r, w_1, w_2)$, where $v_1, \ldots, v_r$ are nodes at level $h+1$. This case is illustrated in Figure 4, where $w_1 = v_6$. Since $j+1$ goes up and only k goes down, and $p_k < p_{j+1}$, T'' has smaller weight than T', a contradiction.

Case 2: There is no leaf at level $\leq h$ in the segment. Consider the cut $(i, i+1, v_1, \ldots, v_r, w_1, w_2)$ where $v_1, \ldots, v_r$ are nodes at level $h+1$. Let $T'' \in \mathrm{OPT}(i)$ be the tree obtained from T by applying *LeftShift*$(i, i+1, v_1, \ldots, v_r, w_1, w_2)$. Only i and $i+1$ go down, and $j+1$ and all the leaves of w_1 go up. Since $p_i + p_{i+1} \leq p_{j+1} < p_{w_1} + p_{j+1}$, T'' has smaller weight than T', a contradiction. This case is illustrated in Figure 5, where $w_1 = v_6$. This completes the proof of Claim B.

Claim C: : If $i > 1$, then $level_{T'}\mathrm{LCA}_{T'}(i-1, i+2) \leq h'$.
Proof of Claim C. The proof is by contradiction. Suppose the claim is false. This implies that there is a node w which is the ancestor of $i-1$ and $i+2$ at level $h'+1$. Let w_1 and w_2 be the sons of w. Let u be the parent of i and $i+1$. Let k be the leftmost node in the cut $(i+3, \ldots j, u)$ which is at level h'. Let $T'' \in \mathrm{OPT}_{moved}(i)$ be the tree obtained from T' by applying *RightShift*$(w_1, w_2, v_1, \ldots, v_r, k)$, where $v_1, \ldots, v_r$ are nodes at level $h'+1$. Since $i-1$ and $i+2$ go up and only k goes down, and $p_k \leq p_i + p_{i+1} < p_{i-1} + p_{i+2}$, T'' has smaller weight than T', a contradiction.

By Claims A, B, and C, the corresponding segments are well shaped in T and T'. This completes the proof in the case that all weights are positive.

We now consider, for completeness, the case where weights may be zero. If $p_i = p_{i+1} = 0$, then $j = i+1$, and the result is trivial. If $p_i + p_{i+1} > 0$, the proof is valid except for one problem, namely that in the proof of Claim B, we must consider the possibility that $p_{w_1} = 0$. Then $r > 0$, since otherwise $p_{i+2} = 0$, contradicting the fact that $(i, i+1)$ is an lmp. Let k be the rightmost leaf of v_r, and let $T'' \in \mathrm{OPT}(i)$ be the tree obtained from T by applying *LeftShift*(k, w_1, w_2). Since only k and leaves of zero weight go down, while $j+1$ goes up, and since $p_k < p_i + p_{i+1} \leq p_{j+1}$, T'' has smaller weight than T', a contradiction. ∎

4 Correctness of the Hu-Tucker algorithm: a simulation

The main idea of the Hu-Tucker algorithm is similar to that of the GW algorithm: combine two items which are very close and whose total weight is small. These items form an active pair which is later combined. However, now, a single item v representing the *combined* active pair is not moved. Instead, v becomes *transparent*. The original items are *opaque*. The algorithm keeps a *working sequence* of names of items, together with their types (opaque or transparent) and weights. A pair of items (u, w) is said to be *compatible* if they are visible to each other, *i.e.*, there is no opaque item between u and w in the current working sequence π. Denote by $pos(u)$ the position of the item u in the leaf-sequence (from left to right).

Definition of a minimal compatible pair. A pair (u, w) of compatible items is a *minimal compatible pair* (*mcp*, for short) if the total weight of (u, w) is minimal. If there are several pairs (u, w) with the same minimal total weight, the pair $(pos(u), pos(w))$ is defined to be the lexically smallest one. The last condition is called the *tie-breaking rule* of the Hu-Tucker algorithm.

Description of the HT **algorithm.** The HT algorithm works in the almost same way as the GW algorithm. Let π be the working sequence, which is initially the original list of items. In the statement $(*)$ in GW, we replace *locally minimal pair* by *minimal compatible pair* and the operation *Move* in the statement $(\#)$ by the statement "make v transparent."

Fix an input sequence of items of length n. Henceforth in this section we assume that there are no ties. This means that no two items in the working sequence π ever have the same weights. The case of ties will be handled in Section 5.

Denote by GW$'$ the *deterministic version* of the GW algorithm in which we choose each time the *globally minimal* lmp, which we call the *gmp.* instead of an arbitrary lmp. Observe that such a pair is not necessarily the rightmost locally minimal pair. This is one of the reasons why we