Minimax Trees in Linear Time
with Applications

Paweł Gawrychowski[1] and Travis Gagie[2],[*]

[1] Institute of Computer Science
University of Wroclaw, Poland
gawry1@gmail.com

[2] Research Group in Genome Informatics
Bielefeld University, Germany
travis.gagie@gmail.com

Abstract. A minimax tree is similar to a Huffman tree except that, instead of minimizing the weighted average of the leaves' depths, it minimizes the maximum of any leaf's weight plus its depth. Golumbic (1976) introduced minimax trees and gave a Huffman-like, $\mathcal{O}(n \log n)$-time algorithm for building them. Drmota and Szpankowski (2002) gave another $\mathcal{O}(n \log n)$-time algorithm, which takes linear time when the weights are already sorted by their fractional parts. In this paper we give the first linear-time algorithm for building minimax trees for unsorted real weights.

1 Introduction

In a minimax tree for a multiset $W = \{w_1, \ldots, w_n\}$ of weights, each leaf has a weight w_i, each internal node has weight equal to the maximum of its children's weights plus 1, and the weight of the root is as small as possible. In other words, if ℓ_i is the depth of the leaf with weight w_i, then $\max_i\{w_i + \ell_i\}$ is minimized. Golumbic [21] showed that if we modify Huffman's algorithm [24] to repeatedly replace the two nodes with smallest weights by a node whose weight is equal to their maximum plus 1, instead of their sum, then it builds a minimax tree instead of a Huffman tree. Like Huffman's algorithm, Golumbic's algorithm takes $\mathcal{O}(n \log n)$ time and can build trees of any degree. Golumbic, Parker [34] and Hoover, Klawe and Pippenger [22] showed how to use Golumbic's algorithm to restrict circuits' fan-in and fan-out without greatly increasing their sizes or depths. While studying prefix codes with minimum maximum pointwise redundancy, Drmota and Szpankowski [10, 11] independently introduced minimax trees as code-trees for generalized Shannon codes [36] and gave another $\mathcal{O}(n \log n)$-time algorithm for building them, which takes linear time when the

---

weights are already sorted by their fractional parts. To see why the two problems are related, consider that, if $P = p_1, \ldots, p_n$ is a probability distribution and each $w_i = \log_2 p_i$, then a minimax tree for W is the code-tree for a prefix code with minimum maximum pointwise redundancy with respect to P. By analyzing their algorithm, Drmota and Szpankowski proved bounds on the redundancy of arithmetic coding, which Baer [3] recently improved by analyzing Golumbic's algorithm. In this paper we show how Drmota and Szpankowski's algorithm can be made to run in linear time on a word RAM when each index and weight fits in $\mathcal{O}(1)$ words. Thus, we obtain the first linear-time algorithm for building minimax trees for unsorted real weights.

Between Golumbic's article and Drmota and Szpankowski's, there seems to have been little research on building minimax trees. Several important papers were published, however, on the related problem of building alphabetic minimax trees, in which the leaves' weights must be in a given order from left to right. Hu, Kleitman and Tamaki [23] gave the first $\mathcal{O}(n \log n)$-time algorithm for building alphabetic minimax trees for real weights. Kirkpatrick and Klawe [27] and Coppersmith, Klawe and Pippenger [6] gave an algorithm (or, more precisely, two algorithms that are equivalent when the trees are to be binary) that builds alphabetic minimax trees for integer weights in linear time, and showed how to use it to restrict circuits' fan-in and fan-out without greatly increasing their sizes or depths and without changing the numbers of edge crossings (and, thus, preserving planarity). Kirkpatrick and Klawe also showed how to combine their algorithm with binary search in order to build alphabet minimax trees for real weights in $\mathcal{O}(n \log n)$ time. We note that, if their algorithm for integer weights is viewed as an alphabetic analogue of the Kraft Inequality [31] — as it was by Yeung [38] and Nakatsu [33], who independently rediscovered it — then their algorithm for real weights is an alphabetic analogue of Drmota and Szpankowski's. Kirkpatrick and Przytycka [28] gave an $\mathcal{O}(\log n)$-time, $\mathcal{O}(n/\log n)$-processor algorithm for integer weights in the CREW PRAM model. Finally, Evans and Kirkpatrick [12] showed how a generalization of Kirkpatrick and Klawe's algorithm can be used to restructure binary search trees.

We became interested in minimax trees while studying adaptive prefix coding. In a previous paper [16] (see also [17, 25]) we noted that minimax trees built with Golumbic's algorithm have the same Sibling Property [13, 18] as Huffman trees, and turned the Faller-Gallager-Knuth algorithm [30] for adaptive Huffman coding into an algorithm for adaptive Shannon coding. Intriguingly, although static Huffman coding is optimal and static Shannon coding is not, adaptive Shannon coding has a better worst-case bound than adaptive Huffman coding does. In another previous paper [15] we used a data structure due to Kirkpatrick and Przytycka and a technique for generalized selection due to Klawe and Mumey [29], to make Kirkpatrick and Klawe's algorithm for real weights run in $\mathcal{O}\bigl(n \min(\log n, d \log \log n)\bigr)$ time, where d is the number of distinct values $\lceil w_i \rceil$. In that paper we conjectured that a similar modification could make Drmota and Szpankowski's algorithm run in linear time on unsorted real weights, and in this paper we prove that conjecture.

In Section 2 we consider the preliminary problem of building minimax trees for unsorted integer weights. Notice that, as such weights have no fractional parts, Drmota and Szpankowski's algorithm takes linear time for this problem. However, there are two difficulties when using their algorithm: first, because they considered the weights to be logarithms, they did not address some questions of precision that arise when the weights are large; second, because they were mostly interested in analysis, they were satisfied with computing the depths of minimax trees' leaves in linear time, rather than building the trees themselves. We give two new linear-time algorithms for unsorted integer weights that can handle large weights — i.e., polynomial in n, so that each fits in a constant number of machine words — and that actually build the minimax trees. In Section 3 we present our main result, a linear-time algorithm for building minimax trees for unsorted real weights. Our algorithm is based on Drmota and Szpankowski's but, whereas theirs uses sorting and binary search, ours uses generalized selection, as well as a new data structure to test the Kraft Inequality. In Section 4 we discuss how our algorithms from Sections 2 and 3 can be applied to problems in, e.g., data compression, group testing and circuit design. In Section 5 we briefly discuss two possible directions for future work. Our results generalize to higher degrees and larger code alphabets but, for the sake of simplicity, in the current version of this paper we consider only binary trees and alphabets; by log we always mean $\log_2$.

2   Minimax Trees for Integer Weights

In this section we give two $\mathcal{O}(n)$-time algorithms for building a minimax tree for a multiset of integer weights, both based on the following lemma (which we note applies to any weights, not only integers) and corollary. We write $M(W)$ to denote the weight of the root of a minimax tree for W.

Lemma 1. *If $W = \{w_1, \ldots, w_n\}$ is a multiset of weights and*

$$W' = \left\{ \max\left(w_1, \max_i\{w_i\} - n + 1\right), \ldots, \max\left(w_n, \max_i\{w_i\} - n + 1\right) \right\},$$

then $M(W') = M(W)$. Moreover, any minimax tree for W' becomes a minimax tree for W when we replace the leaves' weights equal to $\max_i\{w_i\} - n + 1$ by the weights in W less than or equal to $\max_i\{w_i\} - n + 1$, in any order.

Proof. Consider a minimax tree T for W. Without loss of generality, we can assume T is strictly binary — i.e., that every internal node has exactly two children — and, therefore, that it has height at most $n - 1$. If $n = 1$, then $W = w_1 = \max_i\{w_i\} - n + 1$. Otherwise, all the leaves have depth at least 1, so $M(W) \geq \max_i\{w_i\} + 1$. Consider any leaf (if one exists) with weight less than $\max_i\{w_i\} - n + 1$ and depth ℓ. Since $\max_i\{w_i\} - n + 1 + \ell \leq \max_i\{w_i\} < M(W)$, increasing that leaf's weight to $\max_i\{w_i\} - n + 1$ and updating its ancestors' weights, does not change the weight $M(W)$ of the root. It follows that $M(W') = M(W)$.

Now consider a minimax tree T' for W'. If we replace the leaves' weights equal to $\max_i\{w_i\} - n + 1$ by the weights in W less than or equal to $\max_i\{w_i\} - n + 1$ and update all the nodes' weights, then the weight $M(W')$ of the root cannot increase nor, by definition, decrease to less than $M(W)$. Since $M(W') = M(W)$, it follows that the re-weighted tree is a minimax tree for W.   □

Corollary 1. *When all the weights in W are integers, we can sort W' in $\mathcal{O}(n)$ time.*

Proof. When all the weights in W at least $\max_i\{w_i\} - n + 1$ are integers, all the weights in W' are integers in the interval $\left[\max_i\{w_i\} - n + 1, \max_i\{w_i\}\right]$. Since this interval has length $n - 1$, we can sort W' in $\mathcal{O}(n)$ time using either direct addressing, which takes $\mathcal{O}(n)$ extra space, or radix sort, which takes no extra space [14].   □

For our first algorithm, we build and sort W'; build a minimax tree for W' using an implementation of Golumbic's algorithm that takes $\mathcal{O}(n)$ time when the weights are already sorted; and replace the leaves' weights equal to $\max_i\{w_i\} - n + 1$ by the weights in W less than or equal to $\max_i\{w_i\} - n + 1$. We note that Van Leeuwen [37] showed how to implement Huffman's algorithm to take $\mathcal{O}(n)$ time when the weights are already sorted. We could implement Golumbic's algorithm analogously, but we think the implementation below is simpler.

Lemma 2. *Golumbic's algorithm can be implemented to take $\mathcal{O}(n)$ time when the weights are already sorted.*

Proof. We start with the weights stored in a linked list in nondecreasing order, and set a pointer to the head of the list. We then repeat the following procedure until there is only one node left in the list, which is the root of a minimax tree for the given weights: we move the pointer along the list to the last weight less than or equal to the maximum of the first two weights plus 1; remove the first two nodes from the list; make those nodes the children of a new node with weight equal to the maximum of their weights plus one; and insert the new node immediately to the right of the pointer. Notice we remove two nodes for each one we insert, so the total number of nodes is $2n - 1$. Therefore, since the pointer passes over each node once, this implementation takes $\mathcal{O}(n)$ time.   □

Building and sorting W' takes $\mathcal{O}(n)$ time, by Corollary 1; building a minimax tree for W' takes $\mathcal{O}(n)$ time, by Lemma 2; replacing the leaves' weights equal to $\max_i\{w_i\} - n + 1$ by the weights in W less than or equal to $\max_i\{w_i\} - n + 1$ takes $\mathcal{O}(n)$ time, because it can be done in any order. By Lemma 1, the resulting tree is a minimax tree for W.

Theorem 1. *Given a multiset W of n integer weights, we can build a minimax tree for W in $\mathcal{O}(n)$ time.*

Our second algorithm differs in its second step: instead of using Golumbic's algorithm to build a minimax tree for W', we use Kirkpatrick and Klawe's $\mathcal{O}(n)$-time algorithm for integer weights to build an alphabetic minimax tree for the

sequence V consisting of the weights in W' in non-increasing order. The algorithm's correctness follows from the Kraft Inequality.

Theorem 2 (Kraft, 1949). *If there exists a binary tree whose leaves have depths $\ell_1, \ldots, \ell_n$, then $\sum_i 1/2^{\ell_i} \leq 1$. Conversely, if $\sum_i 1/2^{\ell_i} \leq 1$ and $\ell_1 \leq \cdots \leq \ell_n$, then there exists an ordered binary tree whose leaves, from left to right, have depths $\ell_1, \ldots, \ell_n$.*

By the latter part of Theorem 2 and a standard exchange argument — i.e., if a minimax tree contains two leaves such that the deeper one has a higher weight than the shallower one, then we can swap their weights — there exists a minimax tree for W' in which the leaves' weights are non-increasing from left to right. Therefore, by definition, any alphabetic minimax tree for V is a minimax tree for W'.

3   Minimax Trees for Real Weights

In this section we give the first $\mathcal{O}(n)$-time algorithm for building minimax trees for unsorted real weights. As we noted in the introduction, our algorithm is based on Drmota and Szpankowski's algorithm but avoids sorting, which is the step that determines their algorithm's $\mathcal{O}(n \log n)$ complexity. In addition to yielding an optimal algorithm for an interesting problem with applications in, e.g., data compression, group testing and circuit design — described in Section 4 — we believe the techniques we use in this section may be of independent interest.

To build a prefix code with minimum maximum pointwise redundancy with respect to a given probability distribution $P = p_1, \ldots, p_n$, Drmota and Szpankowski start with a Shannon code for P, in which the codeword for the ith character has length $\lceil \log(1/p_i) \rceil$, for each i; they sort the logarithms by their fractional parts, i.e., $\log(1/p_1) - \lfloor \log(1/p_1) \rfloor, \ldots, \log(1/p_n) - \lfloor \log(1/p_n) \rfloor$; and they find the largest value x such that $\lceil \log(1/p_1) - x \rceil, \ldots, \lceil \log(1/p_n) - x \rceil$ obey the Kraft Inequality. A binary tree with leaves at these depths is the code-tree for a prefix code with minimum maximum pointwise redundancy with respect to P, and a minimax tree for $\{\log p_1, \ldots, \log p_n\}$.

We can use Drmota and Szpankowski's algorithm to build a minimax tree given any multiset of weights because, for any value c, if $W = \{w_1, \ldots, w_n\}$ and $W' = \{w_1 + c, \ldots, w_n + c\}$ then, by definition, $M(W') = M(W) + c$ and any minimax tree for W' becomes a minimax tree for W when we subtract c from each leaf's weight. In particular, if $c = -\log\left(\sum_i 2^{w_i}\right)$ then $\sum_i 2^{w_i+c} = 2^c \sum_i 2^{w_i} = 1$; therefore, $W' = \{\log p_1, \ldots, \log p_n\}$ for some probability distribution $P = p_1, \ldots, p_n$ and we can use Drmota and Szpankowski's algorithm directly to build minimax trees for W' and, thus, for W. Without loss of generality, we henceforth assume the given multiset W of weights is equal to $\{\log p_1, \ldots, \log p_n\}$ for some probability distribution P (so, in particular, each $w_i \leq 0$). We can restate the theorem Drmota and Szpankowski proved to establish the correctness of their algorithm — and which also establishes the correctness of our own — in terms of minimax trees instead of prefix codes, as follows:

Theorem 3 (Drmota and Szpankowski, 2002). *If $W = \{w_1, \ldots, w_n\}$ is a multiset of weights, $X = \{x_1, \ldots, x_n\} = \{|w_1| - \lfloor|w_1|\rfloor, \ldots, |w_n| - \lfloor|w_n|\rfloor\}$ and x_i is the largest element in $X \cup \{0\}$ such that*

$$\sum_{x_j \leq x_i} 1/2^{\lfloor|w_j|\rfloor} + \sum_{x_j > x_i} 1/2^{\lceil|w_j|\rceil} \leq 1,$$

then any minimax tree for $\{-\lfloor|w_j|\rfloor : x_j \leq x_i\} \cup \{-\lceil|w_j|\rceil : x_j > x_i\}$ becomes a minimax tree for W when we replace each leaf's weight $-\lfloor|w_j|\rfloor$ or $-\lceil|w_j|\rceil$ by w_j.

If $x_1 \leq \cdots \leq x_n$ and $x_i > 0$ then, by Theorem 3, i is the largest index such that $\{\lfloor|w_j|\rfloor : x_j \leq x_i\} \cup \{\lceil|w_j|\rceil : x_j > x_i\}$ satisfies the Kraft Inequality. To build a minimax tree for W with Drmota and Szpankowski's algorithm, we compute and sort X; use binary search to find i, in each round testing whether the Kraft Inequality holds; build a minimax tree for $\{-\lfloor|w_1|\rfloor, \ldots, -\lfloor|w_i|\rfloor, -\lceil|w_{i+1}|\rceil, \ldots, -\lceil|w_n|\rceil\}$; and replace each leaf's weight $-\lfloor|w_j|\rfloor$ or $-\lceil|w_j|\rceil$ by w_j. Our version differs in three ways: we use generalized selection instead of sorting and binary search; we use a new data structure to test the Kraft Inequality; and we use either of our algorithms from Section 2 to build the minimax tree for $\{-\lfloor|w_1|\rfloor, \ldots, -\lfloor|w_i|\rfloor, -\lceil|w_{i+1}|\rceil, \ldots, -\lceil|w_n|\rceil\}$. In the remainder of this section we first show how to use generalized selection to find i in $\mathcal{O}(n)$ time, excluding the time needed to test the Kraft Inequality; we then show how to perform all the necessary tests in a total of $\mathcal{O}(n)$ time using our new data structure. Since each of our algorithms from Section 2 takes $\mathcal{O}(n)$ time, it follows that we can build a minimax tree for W in $\mathcal{O}(n)$ time.

To find x_i in $\mathcal{O}(n)$ time with general selection, we start with the multiset $X_1 = X \cup \{0\}$ and repeat the following procedure until we reach the empty set: in the rth round, we use the linear-time selection algorithm due to Blum *et al.* [4] to find the current multiset X_r's median x_m, then test whether

$$\sum_{x_j \leq x_m} 1/2^{\lfloor|w_j|\rfloor} + \sum_{x_j > x_m} 1/2^{\lceil|w_j|\rceil} \leq 1;$$

if so, we remove those elements of X_r that are less than or equal to x_m and recurse on the resulting multiset; if not, we remove those elements of X_r that are greater than or equal to x_m and recurse. The element x_i is the largest median we consider for which the test is positive. Since the size of the multisets decreases by a factor of at least 2 in each round, we use $\mathcal{O}(\log n)$ rounds and we find all the medians in a total of $\mathcal{O}(n)$ time.

By the same arguments we used to prove Lemma 1, we can assume, without loss of generality, that $\lceil|w_j|\rceil \leq n - 1$ for each j. To test the Kraft Inequality, we use a data structure consisting of two n-bit binary fractions, S_1 and S_2, each broken into $(\log n)$-bit blocks and initially set to 0. For $1 \leq k \leq n - 1$, adding $1/2^k$ to either fraction takes $\mathcal{O}(1)$ amortized time, for the same reason that incrementing a binary counter takes $\mathcal{O}(1)$ amortized time (see, e.g., [7, Section 17.3]). Nondestructively testing whether $S_1 + S_2 \leq 1$ takes $\mathcal{O}(n/\log n)$

time, because adding each corresponding pair of blocks takes $\mathcal{O}(1)$ time and, by induction, the number carried from each pair to the next is at most 1; resetting either fraction to 0 takes $\mathcal{O}(1)$ time for each block, i.e., $\mathcal{O}(n/\log n)$ time in total.

Before starting to search for x_i, we set $S_1 = \sum_j 1/2^{\lceil |w_j| \rceil}$ in $\mathcal{O}(n)$ time. Throughout our generalized selection, we maintain the invariant that, at the beginning of the rth round,

$$S_1 = \sum_j 1/2^{\lceil |w_j| \rceil} + \sum_{0 < x_j < \min(X_r)} 1/2^{\lceil |w_j| \rceil}$$

and $S_2 = 0$. In the rth round, we set

$$S_2 = \sum_{\min(X_r) \leq x_j \leq x_m} 1/2^{\lceil |w_j| \rceil}$$

in $\mathcal{O}(|X_r|)$ time. Since

$$S_1 + S_2 = \sum_j 1/2^{\lceil |w_j| \rceil} + \sum_{0 < x_j < \min(X_r)} 1/2^{\lceil |w_j| \rceil} + \sum_{\min(X_r) \leq x_j \leq x_m} 1/2^{\lceil |w_j| \rceil}$$
$$= \sum_{x_j \leq x_m} 1/2^{\lfloor |w_j| \rfloor} + \sum_{x_j > x_m} 1/2^{\lceil |w_j| \rceil},$$

we can test the Kraft Inequality in $\mathcal{O}(n/\log n)$ time by checking whether $S_1 + S_2 \leq 1$. If the test is positive, then we add S_2 to S_1 in $\mathcal{O}(n/\log n)$ time; if the test is negative, then we do not change S_1. In either case, straightforward calculation shows that, afterwards,

$$S_1 = \sum_j 1/2^{\lceil |w_j| \rceil} + \sum_{0 < x_j < \min(X_{r+1})} 1/2^{\lceil |w_j| \rceil}$$

so the first part of our invariant is maintained. Finally, we reset $S_2 = 0$ in $\mathcal{O}(n/\log n)$ time, so the second part of our invariant is maintained. Since $|X_r| = \mathcal{O}(n/2^r)$, the rth round takes a total of $\mathcal{O}(n/2^r + n/\log n)$ time. Since $\sum_{r \geq 1} n/2^r = n$ and we use $\mathcal{O}(\log n)$ rounds, it follows that our whole generalized selection takes $\mathcal{O}(n)$ time. This completes the proof of our main result:

Theorem 4. *Given a multiset W of n real weights, we can build a minimax tree for W in $\mathcal{O}(n)$ time.*

4   Applications

Our results from Sections 2 and 3 allow us to solve the following problems in linear time:

    I. build a prefix code with minimum maximum pointwise redundancy with respect to a given distribution;

  II. given a good estimate of the distribution over an alphabet, build a good prefix code;

 III. given a good estimate of the distribution over a set, design a good group test to find the unique target;

 IV. build a minimax tree for a given multiset of real weights;

  V. build a Shannon code for a given distribution;

 VI. build a tree whose leaves have at most given depths;

VII. restrict a given circuit to have bounded fan-in or fan-out;

VIII. build a minimax tree for a given multiset of integer weights.

The authors cited in the introduction have already shown, however, that Problem I takes $\mathcal{O}(n)$ more time than IV [10, 11, 3], V than VI [10, 11, 16], and VI and VII than VIII [21, 34, 22]. Therefore, in the current version of this paper, we discuss only Problems II, III, IV and VIII. We showed in Sections 2 and 3, respectively, that Problem VIII and IV take $\mathcal{O}(n)$ time. In the remainder of this section we define what we mean by "good" in Problems II and III, and show they take $\mathcal{O}(n)$ more time than IV. (Problems II and III are, in fact, equivalent to each other and to I, and analogous to a problem we considered in our paper [15] on building alphabetic minimax trees.) It follows that all the problems listed above take $\mathcal{O}(n)$ time.

    Suppose we want to build a good prefix code with which to compress a file, but are given only a sample of its characters. Let $P = p_1, \ldots, p_n$ be the normalized distribution of characters in the file, let $Q = q_1, \ldots, q_n$ be the normalized distribution of characters in the sample and suppose our codewords are $C = c_1, \ldots, c_n$. An ideal code for Q assigns the ith character a codeword of length $\log(1/q_i)$ (which may not be an integer), and the average codeword's length using such a code is $H(P) + D(P\|Q)$, where $H(P) = \sum_i p_i \log(1/p_i)$ is the entropy of P and $D(P\|Q) = \sum_i p_i \log(p_i/q_i)$ is the relative entropy between P and Q. The entropy measures our expected surprise at a character drawn uniformly at random from the file, given P; the relative entropy (also known as the informational divergence or Kullback-Leibler pseudo-distance) measures the increase in our expected surprise when we estimate P by Q, and is often used to quantify how well Q approximates P (see, e.g., [8]).

    Consider the best worst-case bound we can achieve, given only Q, on how much the average codeword's length exceeds $H(P) + D(P\|Q)$. A result by Katona and Nemetz [26] implies we do not generally achieve a constant bound on the difference when C is a Huffman code for Q. (Given P, of course, the best bound we could achieve on how much the average codeword's length exceeds $H(P)$, would be the redundancy of a Huffman code for P.) For example, if $q_1, \ldots, q_n$ are proportional to $F_n, \ldots, F_1$, where F_i denotes the ith Fibonacci number (i.e., $F_1 = F_2 = 1$ and $F_i = F_{i-1} + F_{i-2}$ for $i \geq 3$), then the codewords' lengths are $1, \ldots, n-2, n-1, n-1$ in any Huffman code for Q. If p_n is sufficiently close to 1, then

$$H(P) + D(P\|Q) \approx \log(1/q_n) = \log \sum_{i=1}^{n} F_i = n \log \phi + \mathcal{O}(1)$$

but the average codeword's length $\sum_i p_i|c_i| \approx n - 1$, so for large n the difference is about $(1/\log \phi - 1)n \approx 0.44n$, where $\phi \approx 1.62$ is the golden ratio.

As long as $q_i > 0$ whenever $p_i > 0$, the average codeword's length

$$\sum_i p_i|c_i| = \sum_i p_i \left(\log(1/p_i) + \log(p_i/q_i) + \log q_i + |c_i| \right)$$

$$= H(P) + D(P\|Q) + \sum_i p_i(\log q_i + |c_i|)$$

(if $q_i = 0$ but $p_i > 0$ for some i, then $D(P\|Q)$ is infinite). Notice each $|c_i|$ is the length of a branch in the code-tree for C. Therefore, the best bound we can achieve is

$$\min_C \max_P \left\{ \sum_i p_i(\log q_i + |c_i|) \right\}$$

$$= \min_C \max_i \{\log q_i + |c_i|\}$$

$$= M(\log q_1, \ldots, \log q_n),$$

which is less than 1 by inspection of Drmota and Szpankowski's algorithm (see also [8, Theorem 5.4.3] and [11, 35, 3]). Moreover, we achieve this bound when the code-tree for C has the same shape as a minimax tree for $\{\log q_1, \ldots, \log q_n\}$. In other words, Problem II takes $\mathcal{O}(n)$ more time than IV.

Now suppose we want to design a good group test (see, e.g., [1, 2]) to find the unique target in a set, given only an estimate Q — presumably gained from past experience or experimentation — of the probability distribution P according to which the target is chosen. A group test allows us to choose, repeatedly, a subset of the elements and check whether the target is among them. We can represent a group test as a decision tree in which each leaf is labelled with an element and each internal node is labelled with the concatenation of its children's labels. Because such a decision tree can be viewed as the code-tree for a prefix code, and vice versa, the expected number of checks we make exceeds $H(P) + D(P\|Q)$ by as little as possible when the decision tree for our group test has the same shape as a minimax tree for $\{\log q_1, \ldots, \log q_n\}$. In other words, Problem III is equivalent to II and, therefore, also takes $\mathcal{O}(n)$ more time than IV.

5   Future Work

We are currently studying whether either Drmota and Szpankowski's solution to Problem I or our solution to II can give us an intuitive explanation of why adaptive Shannon coding has a better worst-case bound than does adaptive Huffman coding. On the one hand, worst-case bounds (especially for online algorithms; see, e.g., [5]) are often proven by considering a game between the algorithm and an omniscient adversary, and minimizing the maximum pointwise redundancy at each step seems somehow related (more than just by name) to the minimax strategy for the algorithm. On the other hand, adaptive prefix coding can be

viewed as a procedure in which we repeatedly build a prefix code based on a sample — i.e., the characters already encoded.

We are still studying alphabetic minimax trees and have started studying minimax trees with unequal edge costs. We believe this latter variant will prove particularly interesting, for three reasons: first, it is not known how to build efficiently a Huffman tree with unequal edge costs (see, e.g., [19, 20]); second, the earliest efficient approximation algorithm (of which we are aware) for building a Huffman tree with unequal edge costs, was Krause's generalization [32] of Shannon coding; third, there is an analogue of the Kraft Inequality for binary trees with unequal edge costs [9].

Acknowledgments

References

1. R. Ahlswede and I. Wegener. *Search Problems.* Wiley, 1987.
2. M. Aigner. *Combinatorial Search.* Wiley, 1988.
3. M. B. Baer. Tight bounds on minimum maximum pointwise redundancy. In *Proceedings of the International Symposium on Information Theory*, pages 1944–1948, 2008.
4. M. Blum, R. W. Floyd, V. R. Pratt, R. L. Rivest, and R. E. Tarjan. Time bounds for selection. *Journal of Computer and System Sciences*, 7(4):448–461, 1973.
5. A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis.* Cambridge University Press, 1998.
6. D. Coppersmith, M. M. Klawe, and N. Pippenger. Alphabetic minimax trees of degree at most *t*. *SIAM Journal on Computing*, 15(1):189–192, 1986.
7. T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms.* MIT Press and McGraw-Hill, 2nd edition, 2001.
8. T. M. Cover and J. A. Thomas. *Elements of Information Theory.* Wiley, 2nd edition, 2006.
9. R. De Prisco and G. Persiano. Characteristic inequalities for binary trees. *Information Processing Letters*, 53(4):201–207, 1995.
10. M. Drmota and W. Szpankowski. Generalized Shannon code minimizes the maximal redundancy. In *Proceedings of the 5th Latin American Symposium on Theoretical Informatics*, pages 306–318, 2002.
11. M. Drmota and W. Szpankowski. Precise minimax redundancy and regret. *IEEE Transactions on Information Theory*, 50(11):2686–2707, 2004.
12. W. S. Evans and D. G. Kirkpatrick. Restructuring ordered binary trees. *Journal of Algorithms*, 50(2):168–193, 2004.
13. N. Faller. An adaptive system for data compression. In *Record of the 7th Asilomar Conference on Circuits, Systems and Computers*, pages 593–597, 1973.
14. G. Franceschini, S. Muthukrishnan, and M. Pǎtraşcu. Radix sorting with no extra space. In *Proceedings of the 15th European Symposium on Algorithms*, pages 194–205, 2007.
15. T. Gagie. A new algorithm for building alphabetic minimax trees. *Fundamenta Informaticae.* To appear.

16. T. Gagie. Dynamic Shannon coding. In *Proceedings of the 12th European Symposium on Algorithms*, pages 359–370, 2004.

17. T. Gagie. Dynamic Shannon coding. *Information Processing Letters*, 102(2–3):113–117, 2007.

18. R. G. Gallager. Variations on a theme by Huffman. *IEEE Transactions on Information Theory*, 24(6):668–674, 1978.

19. M. J. Golin, C. Kenyon, and N. E. Young. Huffman coding with unequal letter costs. In *Proceedings of the 34th Symposium on Theory of Computing*, pages 785–791, 2002.

20. M. J. Golin and J. Li. More efficient algorithms and analyses for unequal letter cost prefix-free coding. *IEEE Transactions on Information Theory*, 52(8):3412–3424, 2008.

21. M. C. Golumbic. Combinatorial merging. *IEEE Transactions on Computers*, 25(11):1164–1167, 1976.

22. H. J. Hoover, M. M. Klawe, and N. Pippenger. Bounding fan-out in logical networks. *Journal of the ACM*, 31(1):13–18, 1984.

23. T. C. Hu, D. J. Kleitman, and J. Tamaki. Binary trees optimum under various criteria. *SIAM Journal on Applied Mathematics*, 37(2):246–256, 1979.

24. D. A. Huffman. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*, 40:1089–1101, 1952.

25. M. Karpinski and Y. Nekrich. A fast algorithm for adaptive prefix coding. *Algorithmica*, 55(1):29–41, 2009.

26. G. O. H. Katona and T. O. H. Nemetz. Huffman codes and self-information. *IEEE Transactions on Information Theory*, 22(3):337–340, 1976.

27. D. G. Kirkpatrick and M. M. Klawe. Alphabetic minimax trees. *SIAM Journal on Computing*, 14(3):514–526, 1985.

28. D. G. Kirkpatrick and T. M. Przytycka. An optimal parallel minimax tree algorithm. In *Proceedings of the 2nd Symposium on Parallel and Distributed Processing*, pages 293–300, 1990.

29. M. M. Klawe and B. Mumey. Upper and lower bounds on constructing alphabetic binary trees. *SIAM Journal on Discrete Mathematics*, 8(4):638–651, 1995.

30. D. E. Knuth. Dynamic Huffman coding. *Journal of Algorithms*, 6(2):163–180, 1985.

31. L. G. Kraft. A device for quantizing, grouping, and coding amplitude-modulated pulses. MSc thesis, Massachusetts Institute of Technology, 1949.

32. R. M. Krause. Channels which transmit letters of unequal duration. *Information and Control*, 5(1):13–24, 1962.

33. N. Nakatsu. Bounds on the redundancy of binary alphabetical codes. *IEEE Transactions on Information Theory*, 37(4):1225–1229, 1991.

34. D. S. Parker, Jr. Combinatorial merging and Huffman's algorithm. *IEEE Transactions on Computers*, 28(5):365–367, 1979.

35. F. Rezaei and C. D. Charalambous. Robust coding for uncertain sources: a minimax approach. In *Proceedings of the International Symposium on Information Theory*, pages 1539–1543, 2005.

36. C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 623–645, 1948.

37. J. van Leeuwen. On the construction of Huffman trees. In *Proceedings of the 3rd International Colloquium on Automata, Languages and Programming*, pages 382–410, 1976.

38. R. W. Yeung. Alphabetic codes revisited. *IEEE Transactions on Information Theory*, 37(3):564–572, 1991.