

D-ary Bounded-Length Huffman Coding

Michael B. Baer
Electronics for Imaging
303 Velocity Way
Foster City, California 94404 USA
Email: Michael.Baer@efi.com

Abstract—Efficient optimal prefix coding has long been accomplished via the Huffman algorithm. However, there is still room for improvement and exploration regarding variants of the Huffman problem. Length-limited Huffman coding, useful for many practical applications, is one such variant, in which codes are restricted to the set of codes in which none of the n codewords is longer than a given length, l_{\max} . Binary length-limited coding can be done in $O(nl_{\max})$ time and $O(n)$ space via the widely used Package-Merge algorithm. In this paper the Package-Merge approach is generalized without increasing complexity in order to introduce a minimum codeword length, l_{\min} , to allow for objective functions other than the minimization of expected codeword length, and to be applicable to both binary and nonbinary codes; nonbinary codes were previously addressed using a slower dynamic programming approach. These extensions have various applications — including faster decompression — and can be used to solve the problem of finding an optimal code with limited fringe, that is, finding the best code among codes with a maximum difference between the longest and shortest codewords. The previously proposed method for solving this problem was nonpolynomial time, whereas solving this using the novel algorithm requires only $O(n(l_{\max} - l_{\min})^2)$ time and $O(n)$ space.

I. INTRODUCTION

A source emits input symbols drawn from the alphabet $\mathcal{X} = \{1, 2, \dots, n\}$, where n is an integer. Symbol i has probability p_i , thus defining probability vector $\mathbf{p} = (p_1, p_2, \dots, p_n)$. Only possible symbols are considered for coding and these are, without loss of generality, sorted in decreasing order of probability; thus $p_i > 0$ and $p_i \leq p_j$ for every $i > j$ such that $i, j \in \mathcal{X}$. Each input symbol is encoded into a codeword composed of output symbols of the D -ary alphabet $\{0, 1, \dots, D-1\}$. The codeword c_i corresponding to input symbol i has length l_i , thus defining length vector $\mathbf{l} = (l_1, l_2, \dots, l_n)$. The code should be a prefix code, i.e., no codeword c_i should begin with the entirety of another codeword c_j .

For the *bounded-length* coding variant of Huffman coding introduced here, all codewords must have lengths lying in a given interval $[l_{\min}, l_{\max}]$. Consider an application in the problem of designing a data codec which is efficient in terms of both compression ratio and coding speed. Moffat and Turpin proposed a variety of efficient implementations of prefix encoding and decoding in [1], each involving table lookups rather than code trees. They noted that the length of the longest codeword should be limited for computational efficiency's sake. Computational efficiency is also improved by restricting the overall range of codeword lengths, reducing the size of the tables and the expected time of searches required for

decoding. Thus, one might wish to have a minimum codeword size of, say, $l_{\min} = 16$ bytes and a maximum codeword size of $l_{\max} = 32$ bytes ($D = 2$). If expected codeword length for an optimal code found under these restrictions is too long, l_{\min} can be reduced and the algorithm rerun until the proper trade-off between coding speed and compression ratio is found.

A similar problem is one of determining opcodes of a microprocessor designed to use variable-length opcodes, each a certain number of bytes ($D = 256$) with a lower limit and an upper limit to size, e.g., a restriction to opcodes being 16, 24, or 32 bits long ($l_{\min} = 2, l_{\max} = 4$). This problem clearly falls within the context considered here, as does the problem of assigning video recorder scheduling codes; these human-readable decimal codes ($D = 10$) also have bounds on their size, such as $l_{\min} = 3$ and $l_{\max} = 8$.

Other problems of interest have $l_{\min} = 0$ and are thus length limited but have no practical lower bound on length [2, p. 396]. Yet other problems have not fixed bounds but a constraint on the difference between minimum and maximum codeword length, a quantity referred to as fringe [3, p. 121]. As previously noted, large fringe has a negative effect of the speed of a decoder.

If we either do not require a minimum or do not require a maximum, it is easy to find values for l_{\min} or l_{\max} that do not limit the problem. As mentioned, setting $l_{\min} = 0$ results in a trivial minimum, as does $l_{\min} = 1$. Similarly, setting $l_{\max} = n$ or using the hard upper bound $l_{\max} = \lceil (n-1)/(D-1) \rceil$ results in a trivial maximum value.

If both minimum and maximum values are trivial, Huffman coding [4] yields a prefix code minimizing expected codeword length $\sum_i p_i l_i$. The conditions necessary and sufficient for the existence of a prefix code with length vector \mathbf{l} are the integer constraint, $l_i \in \mathbb{Z}_+$, and the Kraft (McMillan) inequality [5],

$$\kappa(\mathbf{l}) \triangleq \sum_{i=1}^n D^{-l_i} \leq 1. \quad (1)$$

Finding values for \mathbf{l} is sufficient to find a corresponding code.

It is not always obvious that we should minimize the expected number of questions $\sum_i p_i l_i$ (or, equivalently, the expected number of questions in excess of the first l_{\min} , $\sum_i p_i (l_i - l_{\min})^+$, where x^+ is x if x is positive, 0 otherwise). We generalize and investigate how to minimize the value

$$\sum_{i=1}^n p_i \varphi(l_i - l_{\min}) \quad (2)$$

under the above constraints for any *penalty function* $\varphi(\cdot)$ convex and increasing on \mathbb{R}_+ . Such an additive measurement of cost is called a *quasiarithmetic penalty*, in this case a convex quasiarithmetic penalty.

One such function φ is $\varphi(\delta) = (\delta + l_{\min})^2$, a quadratic value useful in optimizing a communications delay problem [6]. Another function, $\varphi(\delta) = D^{t(\delta + l_{\min})}$ for $t > 0$, can be used to minimize the probability of buffer overflow in a queueing system [7].

Mathematically stating the bounded-length problem,

$$\begin{array}{ll} \text{Given} & \mathbf{p} = (p_1, \dots, p_n), p_i > 0; \\ & D \in \{2, 3, \dots\}; \\ & \text{convex, monotonically increasing} \\ & \varphi : \mathbb{R}_+ \rightarrow \mathbb{R}_+ \\ \text{Minimize } \{l\} & \sum_i p_i \varphi(l_i - l_{\min}) \\ \text{subject to} & \sum_i D^{-l_i} \leq 1; \\ & l_i \in \{l_{\min}, l_{\min} + 1, \dots, l_{\max}\}. \end{array}$$

Note that we need not assume that probabilities p_i sum to 1; they could instead be arbitrary positive weights.

Given a finite n -symbol input alphabet with an associated probability vector \mathbf{p} , a D -symbol output alphabet with codewords of lengths $[l_{\min}, l_{\max}]$ allowed, and a constant-time calculable penalty function φ , we describe an $O(n(l_{\max} - l_{\min}))$ -time $O(n)$ -space algorithm for constructing a φ -optimal code. In Section II, we present a brief review of the relevant literature before extending to D -ary codes a notation first presented in [6]. This notation aids in solving the problem in question by reformulating it as an instance of the D -ary Coin Collector's problem, presented in the section as an extension of the original (binary) Coin Collector's problem [8]. An extension of the Package-Merge algorithm solves this problem; we introduce the reduction and resulting algorithm in Section III. An application to a previously proposed problem involving tree fringe is discussed in Section IV.

II. PRELIMINARIES

Reviewing how the problem in question differs from binary Huffman coding:

- 1) It can be nonbinary, a case considered by Huffman in his original paper [4];
- 2) There is a maximum codeword length, a restriction previously considered, e.g., [9] in $O(n^3 l_{\max} \log D)$ time [10] and $O(n^2 \log D)$ space, but solved efficiently only for binary coding, e.g., [8] in $O(n l_{\max})$ time $O(n)$ space and most efficiently in [11];
- 3) There is a minimum codeword length, a novel restriction;
- 4) The penalty can be nonlinear, a modification previously considered, but only for binary coding, e.g., [12].

The minimum size constraint on codeword length requires a relatively simple change of solution range to [8]. The nonbinary coding generalization is a bit more involved; it requires first modifying the Package-Merge algorithm to allow for an arbitrary numerical base (binary, ternary, etc.), then modifying the coding problem to allow for a provable reduction to the

modified Package-Merge algorithm. The $O(n(l_{\max} - l_{\min}))$ -time $O(n)$ -space algorithm minimizes *height* (that is, maximum codeword length) among optimal codes (if multiple optimal codes exist).

Before presenting an algorithm for optimizing the above problem, we introduce a notation for codes that generalizes one first presented in [6] and modified in [12].

The key idea: Each node (i, l) represents both the share of the penalty (2) (*weight*) and the (scaled) share of the Kraft sum (1) (*width*) assumed for the l th bit of the i th codeword. By showing that total weight is an increasing function of the penalty and that there is a one-to-one correspondence between an optimal code and a corresponding optimal nodeset, we reduce the problem to an efficiently solvable problem, the Coin Collector's problem.

In order to do this, we first need to make a modification to the problem analogous to one Huffman made in his original nonbinary solution. We must in some cases add a “dummy” input or “dummy” inputs of probability $p_i = 0$ to the probability vector to assure that the optimal code has the Kraft inequality satisfied with equality, an assumption underlying both the Huffman algorithm and ours. If we use the minimum number of dummy inputs needed to make $n \bmod (D - 1) \equiv 1$, we can assume without loss of generality that $\kappa(l) = 1$. With this modification, we present *nodeset* notation:

Definition 1: A *node* is an ordered pair of integers (i, l) such that $i \in \{1, \dots, n\}$ and $l \in \{l_{\min} + 1, \dots, l_{\max}\}$. Call the set of all possible nodes I . This set can be arranged in an $n \times (l_{\max} - l_{\min})$ grid, e.g., Fig. 1. The set of nodes, or *nodeset*, corresponding to input symbol i (assigned codeword c_i with length l_i) is the set of the first $l_i - l_{\min}$ nodes of column i , that is, $\eta(i) \triangleq \{(j, l) \mid j = i, l \in \{l_{\min} + 1, \dots, l_i\}\}$. The nodeset corresponding to length vector \mathbf{l} is $\eta(\mathbf{l}) \triangleq \bigcup_i \eta(i)$; this corresponds to a set of n codewords, a code. Thus, in Fig. 1, the dashed line surrounds a nodeset corresponding to $\mathbf{l} = (1, 2, 2, 2, 2, 2)$. We say a node (i, l) has *width* $\rho(i, l) \triangleq D^{-l}$ and *weight* $\mu(i, l) \triangleq p_i \varphi(l - l_{\min}) - p_i \varphi(l - l_{\min} - 1)$, as shown in the example in Fig. 1. Note that if $\varphi(l) = l$, $\mu(i, l) = p_i$.

Given valid nodeset $N \subseteq I$, it is straightforward to find the corresponding length vector and, if it satisfies the Kraft inequality, a code.

We find an optimal nodeset using the D -ary Coin Collector's problem. Let $D^{\mathbb{Z}}$ denote the set of all integer powers of a fixed integer $D > 1$. The Coin Collector's problem of size m considers “coins” indexed by $i \in \{1, 2, \dots, m\}$. Each coin has a width, $\rho_i \in D^{\mathbb{Z}}$; one can think of width as coin face value, e.g., $\rho_i = 0.25 = 2^{-2}$ for a quarter dollar (25 cents). Each coin also has a weight, $\mu_i \in \mathbb{R}$. The final problem parameter is total width, denoted ρ_{tot} . The problem is then:

$$\begin{array}{ll} \text{Minimize } \{B \subseteq \{1, \dots, m\}\} & \sum_{i \in B} \mu_i \\ \text{subject to} & \sum_{i \in B} \rho_i = \rho_{\text{tot}} \\ \text{where} & m \in \mathbb{Z}_+, \mu_i \in \mathbb{R} \\ & \rho_i \in D^{\mathbb{Z}}, \rho_{\text{tot}} \in \mathbb{R}_+. \end{array} \quad (3)$$

We thus wish to choose coins with total width ρ_{tot} such that their total weight is as small as possible. This problem has a linear-time solution given sorted inputs; this solution was found for $D = 2$ in [8] and is found for $D > 2$ here.

Let $i \in \{1, \dots, m\}$ denote both the index of a coin and the coin itself, and let \mathcal{I} represent the m items along with their weights and widths. The optimal solution, a function of total width ρ_{tot} and items \mathcal{I} , is denoted $\text{CC}(\mathcal{I}, \rho_{\text{tot}})$ (the optimal coin collection for \mathcal{I} and ρ_{tot}). Note that, due to ties, this need not be a unique solution, but the algorithm proposed here is deterministic; that is, it finds one specific solution, much like bottom-merge Huffman coding [13] or the corresponding length-limited problem [12], [14]

Because we only consider cases in which a solution exists, $\rho_{\text{tot}} = \omega \rho_{\text{pow}}$ for some $\rho_{\text{pow}} \in D^{\mathbb{Z}}$ and $\omega \in \mathbb{Z}_+$. Here, assuming $\rho_{\text{tot}} > 0$, ρ_{pow} and ω are the unique pair of a power of D and an integer that is not a multiple of D , respectively, which, multiplied, form ρ_{tot} . If $\rho_{\text{tot}} = 0$, ω and ρ_{pow} are not used. Note that ρ_{pow} need not be an integer.

Algorithm variables

At any point in the algorithm, given nontrivial \mathcal{I} and ρ_{tot} , we use the following definitions:

Remainder	$\rho_{\text{pow}} \triangleq$	the unique $x \in D^{\mathbb{Z}}$ such that $\frac{\rho_{\text{tot}}}{x} \in \mathbb{Z} \setminus D\mathbb{Z}$
Minimum width	$\rho^* \triangleq$	$\min_{i \in \mathcal{I}} \rho_i \ (\in D^{\mathbb{Z}})$
Small width set	$\mathcal{I}^* \triangleq$	$\{i \mid \rho_i = \rho^*\} \ (\neq \emptyset)$
“First” item	$i^* \triangleq$	$\arg \min_{i \in \mathcal{I}^*} \mu_i$ (ties broken w/highest index)
“First” package	$\mathcal{P}^* \triangleq$	$\begin{cases} \mathcal{P} \text{ such that} \\ \mathcal{P} = D, \\ \mathcal{P} \subseteq \mathcal{I}^*, \\ \mathcal{P} \preceq \mathcal{I}^* \setminus \mathcal{P}, \quad \mathcal{I}^* \geq D \\ \emptyset, \quad \mathcal{I}^* < D \end{cases}$ (ties broken w/highest indices)

where $D\mathbb{Z}$ denotes integer multiples of D and $\mathcal{P} \preceq \mathcal{I}^* \setminus \mathcal{P}$ denotes that, for all $i \in \mathcal{P}$ and $j \in \mathcal{I}^* \setminus \mathcal{P}$, $\mu_i \leq \mu_j$. Then the following is a recursive description of the algorithm:

Recursive D -ary Package-Merge Procedure

Basis. $\rho_{\text{tot}} = 0$: $\text{CC}(\mathcal{I}, \rho_{\text{tot}}) = \emptyset$.

Case 1. $\rho^* = \rho_{\text{pow}}$ and $\mathcal{I} \neq \emptyset$: $\text{CC}(\mathcal{I}, \rho_{\text{tot}}) = \text{CC}(\mathcal{I} \setminus \{i^*\}, \rho_{\text{tot}} - \rho^*) \cup \{i^*\}$.

Case 2a. $\rho^* < \rho_{\text{pow}}$, $\mathcal{I} \neq \emptyset$, and $|\mathcal{I}^*| < D$: $\text{CC}(\mathcal{I}, \rho_{\text{tot}}) = \text{CC}(\mathcal{I} \setminus \mathcal{I}^*, \rho_{\text{tot}})$.

Case 2b. $\rho^* < \rho_{\text{pow}}$, $\mathcal{I} \neq \emptyset$, and $|\mathcal{I}^*| \geq D$: Create i' , a new item with weight $\mu_{i'} = \sum_{i \in \mathcal{P}^*} \mu_i$ and width $\rho_{i'} = D\rho^*$. This new item is thus a combined item, or *package*, formed by combining the D least weighted items of width ρ^* . Let $\mathcal{S} = \text{CC}(\mathcal{I} \setminus \mathcal{P}^* \cup \{i'\}, \rho_{\text{tot}})$ (the optimization of the packaged

version, where the package is given a low index so that, if “repackaged,” this occurs after all singular or previously packaged items of identical weight and width). If $i' \in \mathcal{S}$, then $\text{CC}(\mathcal{I}, \rho_{\text{tot}}) = \mathcal{S} \setminus \{i'\} \cup \mathcal{P}^*$; otherwise, $\text{CC}(\mathcal{I}, \rho_{\text{tot}}) = \mathcal{S}$.

Theorem 1: If an optimal solution to the Coin Collector’s problem exists, the above recursive (Package-Merge) algorithm will terminate with an optimal solution.

Proof: Using induction on the number of input items, while the basis is trivially correct, each inductive case reduces the number of items by at least one. The inductive hypothesis on $\rho_{\text{tot}} \geq 0$ and $\mathcal{I} \neq \emptyset$ is that the algorithm is correct for any problem instance with fewer input items than instance $(\mathcal{I}, \rho_{\text{tot}})$.

If $\rho^* > \rho_{\text{pow}} > 0$, or if $\mathcal{I} = \emptyset$ and $\rho_{\text{tot}} \neq 0$, then there is no solution to the problem, contrary to our assumption. Thus all feasible cases are covered by those given in the procedure. Case 1 indicates that the solution must contain at least one element (item or package) of width ρ^* . These must include the minimum weight item in \mathcal{I}^* , since otherwise we could substitute one of the items with this “first” item and achieve improvement. Case 2 indicates that the solution must contain a number of elements of width ρ^* that is a multiple of D . If this number is 0, none of the items in \mathcal{P}^* are in the solution. If it is not, then they all are. Thus, if $\mathcal{P}^* = \emptyset$, the number is 0, and we have Case 2a. If not, we may “package” the items, considering the replaced package as one item, as in Case 2b. Thus the inductive hypothesis holds. ■

The algorithm can be performed in linear time and space, as with the binary version [8].

III. A GENERAL ALGORITHM

Theorem 2: The solution N of the Package-Merge algorithm for $\mathcal{I} = I$ and

$$\rho_{\text{tot}} = \frac{n - D^{l_{\min}}}{D - 1} D^{-l_{\min}}$$

has a corresponding length vector \mathbf{l}^N such that $N = \eta(\mathbf{l}^N)$ and $\mu(N) = \min_l \sum_i p_i \varphi(l_i - l_{\min}) - \varphi(0) \sum_i p_i$.

A formal proof can be found in the full version at [15]. The idea is to show that, if there is an $(i, l) \in N$ with $l \in [l_{\min} + 2, l_{\max}]$ such that $(i, l - 1) \in I \setminus N$, one can strictly decrease the penalty by substituting item $(i, l - 1)$ for a set of items including (i, l) , showing the suboptimality of N . Conversely, if there is no such (i, l) , optimal N corresponds to an optimal length vector.

Because the Coin Collector’s problem is linear in time and space — same-width inputs are presorted by weight, numerical operations and comparisons are constant time — the overall algorithm finds an optimal code in $O(|\mathcal{I}|) = O(n(l_{\max} - l_{\min}))$ time and space. Space complexity, however, can be lessened. This is because the algorithm output is a monotonic nodeset:

Definition 2: A *monotonic nodeset*, N , is one with the following properties:

$$(i, l) \in N \Rightarrow (i + 1, l) \in N \quad \text{for } i < n \quad (4)$$

$$(i, l) \in N \Rightarrow (i, l - 1) \in N \quad \text{for } l > l_{\min} + 1. \quad (5)$$




Fig. 1. The set of nodes I with widths $\{\rho(i, l)\}$ and weights $\{\mu(i, l)\}$ for $\varphi(\delta) = \delta^2$, $n = 7$, $D = 3$, $l_{\min} = 1$, $l_{\max} = 4$

In other words, a nodeset is monotonic if and only if it corresponds to a length vector l with lengths sorted in increasing order; this definition is equivalent to that given in [8].

While not all optimal codes are monotonic, using the aforementioned tie-breaking techniques, the algorithm always results in a monotonic code, one that has minimum maximum length among all monotonic optimal codes. Examples of monotonic nodesets include the sets of nodes enclosed by dashed lines in Fig. 1 and Fig. 2. In the latter case, $n = 21$, $D = 3$, $l_{\min} = 2$, and $l_{\max} = 8$, so $\rho_{\text{tot}} = 2/3$.

In [8], monotonicity allows trading off a constant factor of time for drastically reduced space complexity for length-limited binary codes. We extend this to the bounded-length problem. Note that the total width of items that are each less than or equal to width ρ is less than $2n\rho$. Thus, when we are processing items and packages of width ρ , fewer than $2n$ packages are kept in memory. The key idea in reducing space complexity is to keep only four attributes of each package in memory instead of the full contents. In this manner, we use $O(n)$ space while retaining enough information to reconstruct the optimal nodeset in algorithmic postprocessing.

Package attributes allow us to divide the problem into two subproblems with total complexity that is at most half that of the original problem. Define

$$l_{\text{mid}} \triangleq \left\lfloor \frac{1}{2}(l_{\max} + l_{\min} + 1) \right\rfloor.$$

For each package S , we retain only the following attributes:

- 1) $\mu(S) \triangleq \sum_{(i,l) \in S} \mu(i, l)$
- 2) $\rho(S) \triangleq \sum_{(i,l) \in S} \rho(i, l)$
- 3) $\nu(S) \triangleq |S \cap I_{\text{mid}}|$
- 4) $\psi(S) \triangleq \sum_{(i,l) \in S \cap I_{\text{hi}}} \rho(i, l)$

where $I_{\text{hi}} \triangleq \{(i, l) \mid l > l_{\text{mid}}\}$ and $I_{\text{mid}} \triangleq \{(i, l) \mid l = l_{\text{mid}}\}$. We also define $I_{\text{lo}} \triangleq \{(i, l) \mid l < l_{\text{mid}}\}$.

With only these parameters, the “first run” of the algorithm takes $O(n)$ space. The output of this run is the package

attributes of the optimal nodeset N . Thus, at the end of this first run, we know the value for $n_\nu \triangleq \nu(N)$, and we can consider N as the disjoint union of four sets, shown in Fig. 2:

- 1) $A = \text{nodes in } N \cap I_{\text{lo}} \text{ with indexes in } [1, n - n_\nu]$,
- 2) $B = \text{nodes in } N \cap I_{\text{lo}} \text{ with indexes in } [n - n_\nu + 1, n]$,
- 3) $\Gamma = \text{nodes in } N \cap I_{\text{mid}}$,
- 4) $\Delta = \text{nodes in } N \cap I_{\text{hi}}$.

Due to the monotonicity of N , it is clear that $B = [n - n_\nu + 1, n] \times [l_{\min} + 1, l_{\text{mid}} - 1]$ and $\Gamma = [n - n_\nu + 1, n] \times \{l_{\text{mid}}\}$. Note then that $\rho(B) = (n_\nu)(D^{-l_{\min}} - D^{1-l_{\text{mid}}})/(D-1)$ and $\rho(\Gamma) = n_\nu D^{-l_{\text{mid}}}$. Thus we need merely to recompute which nodes are in A and in Δ .

Because Δ is a subset of I_{hi} , $\rho(\Delta) = \psi(N)$ and $\rho(A) = \rho(N) - \rho(B) - \rho(\Gamma) - \rho(\Delta)$. Given their respective widths, A is a minimal weight subset of $[1, n - n_\nu] \times [l_{\min} + 1, l_{\text{mid}} - 1]$ and Δ is a minimal weight subset of $[n - n_\nu + 1, n] \times [l_{\text{mid}} + 1, l_{\max}]$. These will be monotonic if the overall nodeset is monotonic. The nodes at each level of A and Δ can thus be found by recursive calls to the algorithm. This approach uses only $O(n)$ space while preserving time complexity as in [8].

There are changes we can make to the algorithm that, for certain inputs, will result in even better performance. For example, if $l_{\max} \approx \log_D n$, then, rather than minimizing the weight of nodes of a certain total width, it is easier to maximize weight and find the complementary set of nodes. Similarly, if most input symbols have one of a handful of probability values, one can consider this and simplify calculations. These and other similar optimizations have been done in the past for the special case $\varphi(\delta) = \delta$, $l_{\min} = 0$, $D = 2$ [16]–[20], though we will not address or extend such improvements here.

Note also that there are cases in which we can find a better upper bound for codeword length than l_{\max} or a better lower bound than l_{\min} . In such cases, complexity is accordingly reduced, and, when l_{\max} is effectively trivial (e.g., $l_{\max} = n - 1$), and the Package-Merge approach can be replaced




Fig. 2. The set of nodes I , an optimal nodeset N , and disjoint subsets A , B , Γ , Δ

by conventional (linear-time) Huffman coding approaches. Likewise, when $\varphi(\delta) = \delta$ and $l_{\max} - l_{\min}$ is not $O(\log n)$, an approach similar to that of [21] as applied in [11] has better asymptotic performance. These alternative approaches are omitted due to space and can be found at [15].

IV. FRINGE-LIMITED PREFIX CODING

An important problem that can be solved with the techniques in this paper is that of finding an optimal code given an upper bound on fringe, the difference between minimum and maximum codeword length; such a problem is proposed in [3, p. 121], where it is suggested that if there are $b - 1$ codes better than the best code with fringe at most d , one can find this b -best code with the $O(bn^3)$ -time algorithm in [22, pp. 890–891], thus solving the fringe-limited problem. However, this presumes we know an upper bound for b before running this algorithm. More importantly, if a probability vector is far from uniform, b can be very large, since the number of viable code trees is $\Theta(1.794 \dots^n)$ [23], [24]. Thus this is a poor approach in general. Instead, we can use the aforementioned algorithms for finding the optimal bounded-length code with codeword lengths restricted to $[l' - d, l']$ for each $l' \in \{\lceil \log_D n \rceil, \lceil \log_D n \rceil + 1, \dots, \lceil \log_D n \rceil + d\}$, keeping the best of these codes; this covers all feasible cases of fringe upper bounded by d . (Here we again assume, without loss of generality, that $n \bmod (D - 1) \equiv 1$.) The overall procedure thus has time complexity $O(nd^2)$ and $O(n)$ space complexity.

ACKNOWLEDGMENTS

The author wishes to thank Zhen Zhang for first bringing a related problem to his attention and David Morgenthaler for constructive discussions on this topic.

REFERENCES

- [1] A. Moffat and A. Turpin, "On the implementation of minimum redundancy prefix codes," *IEEE Trans. Commun.*, vol. 45, no. 10, pp. 1200–1207, Oct. 1997.
- [2] I. H. Witten, A. Moffat, and T. Bell, *Managing Gigabytes*, 2nd ed. San Francisco, CA: Morgan Kaufmann Publishers, 1999.
- [3] J. Abrahams, "Code and parse trees for lossless source encoding," *Communications in Information and Systems*, vol. 1, no. 2, pp. 113–146, Apr. 2001.
- [4] D. A. Huffman, "A method for the construction of minimum-redundancy codes," *Proc. IRE*, vol. 40, no. 9, pp. 1098–1101, Sept. 1952.

- [5] B. McMillan, "Two inequalities implied by unique decipherability," *IRE Trans. Inf. Theory*, vol. IT-2, no. 4, pp. 115–116, Dec. 1956.
- [6] L. L. Larmore, "Minimum delay codes," *SIAM J. Comput.*, vol. 18, no. 1, pp. 82–94, Feb. 1989.
- [7] P. A. Humblet, "Generalization of Huffman coding to minimize the probability of buffer overflow," *IEEE Trans. Inf. Theory*, vol. IT-27, no. 2, pp. 230–232, Mar. 1981.
- [8] L. L. Larmore and D. S. Hirschberg, "A fast algorithm for optimal length-limited Huffman codes," *J. ACM*, vol. 37, no. 2, pp. 464–473, Apr. 1990.
- [9] A. Itai, "Optimal alphabetic trees," *SIAM J. Comput.*, vol. 5, no. 1, pp. 9–18, Mar. 1976.
- [10] L. Gottlieb and D. Wood, "The construction of optimal multiway search trees and the monotonicity principle," *Intern. J. Computer Maths, Section A*, vol. 9, no. 1, pp. 17–24, 1981.
- [11] B. Schieber, "Computing a minimum-weight k -link path in graphs with the concave Monge property," *Journal of Algorithms*, vol. 29, no. 2, pp. 204–222, Nov. 1998.
- [12] M. B. Baer, "Source coding for quasiarithmetic penalties," *IEEE Trans. Inf. Theory*, vol. IT-52, no. 10, pp. 4380–4393, Oct. 2006.
- [13] E. S. Schwartz, "An optimum encoding with minimum longest code and total number of digits," *Inf. Contr.*, vol. 7, no. 1, pp. 37–44, Mar. 1964.
- [14] L. L. Larmore and T. M. Przytycka, "A fast algorithm for optimal height-limited alphabetic binary-trees," *SIAM J. Comput.*, vol. 23, no. 6, pp. 1283–1312, Dec. 1994.
- [15] M. B. Baer, "Twenty (or so) questions: D -ary bounded-length Huffman coding," preprint available from <http://arxiv.org/abs/cs.IT/0602085>.
- [16] J. Katajainen, A. Moffat, and A. Turpin, "A fast and space-economical algorithm for length-limited coding," in *Proc., Int. Symp. on Algorithms and Computation*, Dec. 1995, p. 1221.
- [17] M. Liddell and A. Moffat, "Incremental calculation of optimal length-restricted codes," in *Proc., IEEE Data Compression Conf.*, Apr. 2–4, 2002, pp. 182–191.
- [18] A. Moffat, A. Turpin, and J. Katajainen, "Space-efficient construction of optimal prefix codes," in *Proc., IEEE Data Compression Conf.*, Mar. 28–30, 1995, pp. 192–202.
- [19] A. Turpin and A. Moffat, "Practical length-limited coding for large alphabets," *The Comput. J.*, vol. 38, no. 5, pp. 339–347, 1995.
- [20] —, "Efficient implementation of the package-merge paradigm for generating length-limited codes," in *Proc., Computing: The Australasian Theory Symposium*, Jan. 29–30, 1996, pp. 187–195.
- [21] L. L. Larmore and T. M. Przytycka, "Parallel construction of trees with optimal weighted path length," in *Proc. 3rd Annual Symposium on Parallel Algorithms and Architectures*, 1991, pp. 71–80.
- [22] S. Anily and R. Hassin, "Ranking the best binary trees," *SIAM J. Comput.*, vol. 18, no. 5, pp. 882–892, Oct. 1989.
- [23] J. Komlos, W. Moser, and T. Nemetz, "On the asymptotic number of prefix codes," *Mitteilungen aus dem Mathematischen Seminar Giessen*, Heft 165, Coxeter Festschrift, Teil III, pp. 35–48, 1984.
- [24] P. Flajolet and H. Prodinger, "Level number sequences for trees," *Disc. Math.*, vol. 65, no. 2, pp. 149–156, June 1987.