

An Isabelle/HOL Formalization of the Textbook Proof of Huffman's Algorithm*

Jasmin Christian Blanchette

Institut für Informatik, Technische Universität München, Germany

blanchette@in.tum.de

October 21, 2008

Abstract

Huffman's algorithm is a procedure for constructing a binary tree with minimum weighted path length. This report presents a formal proof of the correctness of Huffman's algorithm written using Isabelle/HOL. Our proof closely follows the sketches found in standard algorithms textbooks, uncovering a few snags in the process. Another distinguishing feature of our formalization is the use of custom induction rules to help Isabelle's automatic tactics, leading to very short proofs for most of the lemmas.

Contents

1	Introduction	2
1.1	Binary Codes	2
1.2	Binary Trees	3
1.3	Huffman's Algorithm	5
1.4	The Textbook Proof	6
1.5	Overview of the Formalization	7
1.6	Overview of Isabelle's HOL Logic	8
1.7	Head of the Theory File	8
2	Definition of Prefix Code Trees and Forests	9
2.1	Tree Datatype	9
2.2	Forest Datatype	9
2.3	Alphabet	9

*This work was supported by the DFG grant NI 491/11-1.

2.4	Consistency	10
2.5	Symbol Depths	11
2.6	Height	12
2.7	Symbol Frequencies	13
2.8	Weight	13
2.9	Cost	14
2.10	Optimality	15
3	Functional Implementation of Huffman's Algorithm	16
3.1	Cached Weight	16
3.2	Tree Union	16
3.3	Ordered Tree Insertion	17
3.4	The Main Algorithm	18
4	Definition of Auxiliary Functions Used in the Proof	18
4.1	Sibling of a Symbol	18
4.2	Leaf Interchange	21
4.3	Symbol Interchange	24
4.4	Four-Way Symbol Interchange	25
4.5	Sibling Merge	27
4.6	Leaf Split	29
4.7	Weight Sort Order	30
4.8	Pair of Minimal Symbols	31
5	Formalization of the Textbook Proof	31
5.1	Four-Way Symbol Interchange Cost Lemma	31
5.2	Leaf Split Optimality Lemma	32
5.3	Leaf Split Commutativity Lemma	34
5.4	Optimality Theorem	36
6	Related Work	37
7	Conclusion	38

1 Introduction

1.1 Binary Codes

Suppose we want to encode strings over a finite source alphabet to sequences of bits. The approach used by ASCII and most other charsets is to map each source symbol to a distinct k -bit codeword, where k is fixed and is typically 8 or 16. To encode a string of symbols, we simply encode each symbol in turn. Decoding involves mapping each k -bit block back to the symbol it represents.

Fixed-length codes are simple and fast, but they generally waste space. If we know the frequency w_a of each source symbol a , we can save some bits by using shorter codewords for the most frequent symbols. We say that a (variable-length) code is *optimum* if it minimizes the sum $\sum_a w_a \delta_a$, where δ_a is the length of the binary codeword for a . Information theory tells us that a code is optimum if for each source symbol c the codeword representing c has length

$$\delta_c = \log_2 \frac{1}{p_c}, \quad \text{where } p_c = \frac{w_c}{\sum_a w_a}.$$

This number is generally not an integer, so we cannot use it directly. Nonetheless, the above criterion is a useful yardstick and paves the way for arithmetic coding [10], a generalization of the method presented here.

As an example, consider the source string 'abacabad'. We have

$$p_a = \frac{1}{2}, \quad p_b = \frac{1}{4}, \quad p_c = \frac{1}{8}, \quad p_d = \frac{1}{8}.$$

The optimum lengths for the binary codewords are all integers, namely

$$\delta_a = 1, \quad \delta_b = 2, \quad \delta_c = 3, \quad \delta_d = 3,$$

and they are realized by the code

$$C_1 = \{a \mapsto 0, b \mapsto 10, c \mapsto 110, d \mapsto 111\}.$$

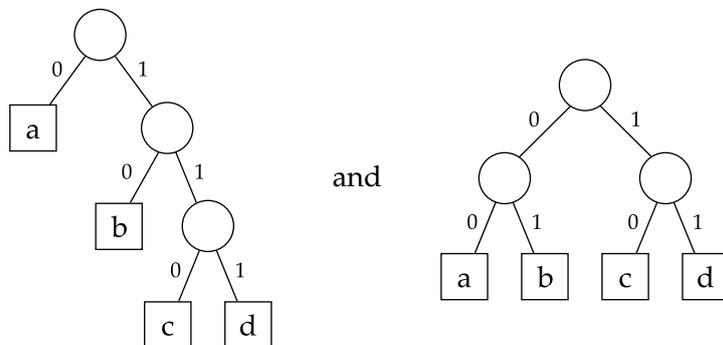
Encoding 'abacabad' produces the 14-bit codeword 01001100100111. The code C_1 is optimum: No code that unambiguously encodes source symbols one at a time could do better than C_1 on the input 'abacabad'. In particular, with a fixed-length code such as

$$C_2 = \{a \mapsto 00, b \mapsto 01, c \mapsto 10, d \mapsto 11\}$$

we need at least 16 bits to encode 'abacabad'.

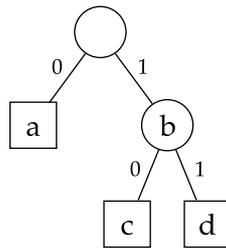
1.2 Binary Trees

Inside a program, binary codes like C_1 and C_2 can be represented by binary trees. For example, the trees



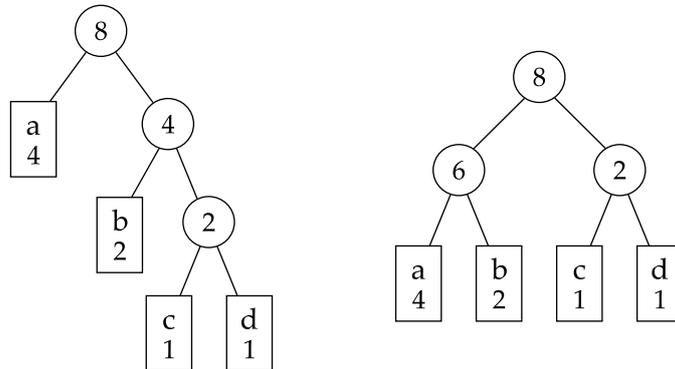
correspond to C_1 and C_2 , respectively. The codeword for a symbol is given along the path from the root to that symbol, with 0 meaning “left child” and 1 meaning “right child”.

To avoid ambiguities, we require that only leaf nodes are labeled with symbols. This ensures that no codeword is a prefix of another, thereby eliminating the source of all ambiguities.¹ Codes that have this property are called *prefix codes*. As an example of a code that doesn't have this property, consider the code associated with the tree



and observe that ‘bbb’, ‘bd’, and ‘db’ all map to the codeword 111.

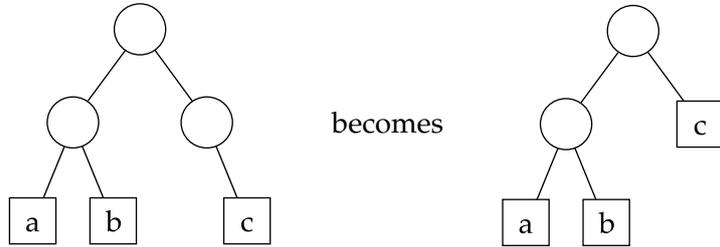
Each node in a code tree is assigned a *weight*. For a leaf node, the weight is the frequency of its symbol; for an inner node, it is the sum of the weights of its subtrees. Code trees can be annotated with their weights:



For our purposes, it is sufficient to consider only full binary trees (trees whose inner nodes all have two children). This is because any inner node with only one

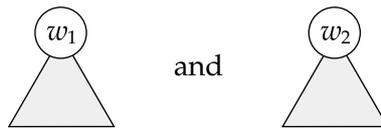
¹Strictly speaking, there is another potential source of ambiguity. If the alphabet consists of a single symbol a , that symbol could be mapped to the empty codeword, and then any string $aa \dots a$ would map to the empty bit sequence, giving the decoder no way to recover the original string's length. This scenario can be ruled out by requiring that the alphabet has cardinality 2 or more.

child can advantageously be eliminated; for example,

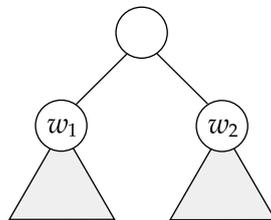


1.3 Huffman's Algorithm

David Huffman [5] discovered a simple algorithm for constructing an optimum prefix code tree for specified symbol frequencies: Create a forest consisting of only leaf nodes, one for each symbol in the alphabet. Then take the two trees



with the lowest weights and replace them with the tree

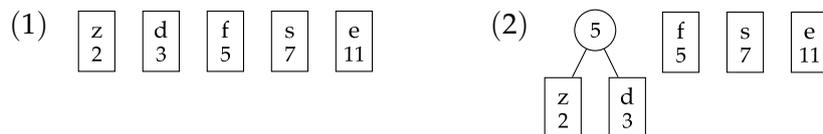


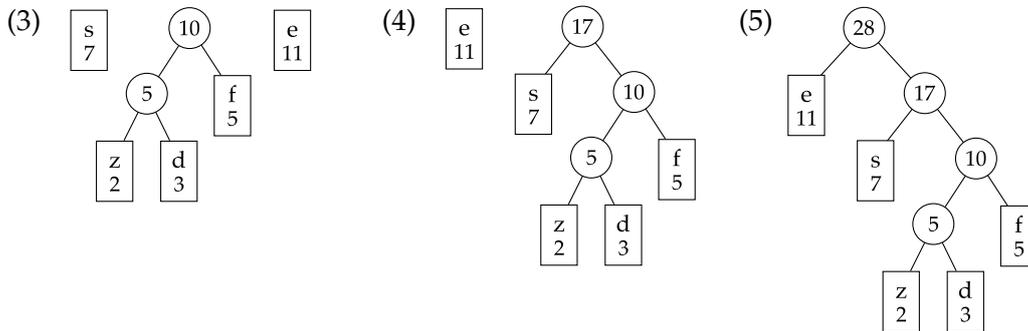
Repeat this process until only one tree is left.

As an illustration, executing the algorithm for the frequencies

$$f_d = 3, f_e = 11, f_f = 5, f_s = 7, f_z = 2$$

gives rise to the following sequence of states:



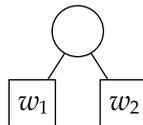


Tree (5) is an optimum tree for the given frequencies.

1.4 The Textbook Proof

Why does the algorithm work? In his article, Huffman gave some motivation but no real proof. For a proof sketch, we turn to Donald Knuth [6, p. 403–404]:

It is not hard to prove that this method does in fact minimize the weighted path length [i.e., $\sum_a w_a \delta_a$], by induction on m . Suppose we have $w_1 \leq w_2 \leq w_3 \leq \dots \leq w_m$, where $m \geq 2$, and suppose that we are given a tree that minimizes the weighted path length. (Such a tree certainly exists, since only finitely many binary trees with m terminal nodes are possible.) Let V be an internal node of maximum distance from the root. If w_1 and w_2 are not the weights already attached to the children of V , we can interchange them with the values that are already there; such an interchange does not increase the weighted path length. Thus there is a tree that minimizes the weighted path length and contains the subtree

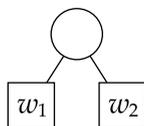


Now it is easy to prove that the weighted path length of such a tree is minimized if and only if the tree with



has minimum path length for the weights $w_1 + w_2, w_3, \dots, w_m$.

There is, however, a small oddity in this proof: It is not clear why we must assert the existence of an optimum tree that contains the subtree



(The formalization works without it.)

Cormen et al. [3, p. 385–391] provide a very similar proof, articulated around the following propositions:

Lemma 16.2

Let C be an alphabet in which each character $c \in C$ has frequency $f[c]$. Let x and y be two characters in C having the lowest frequencies. Then there exists an optimal prefix code for C in which the codewords for x and y have the same length and differ only in the last bit.

Lemma 16.3

Let C be a given alphabet with frequency $f[c]$ defined for each character $c \in C$. Let x and y be two characters in C with minimum frequency. Let C' be the alphabet C with characters x, y removed and (new) character z added, so that $C' = C - \{x, y\} \cup \{z\}$; define f for C' as for C , except that $f[z] = f[x] + f[y]$. Let T' be any tree representing an optimal prefix code for the alphabet C' . Then the tree T , obtained from T' by replacing the leaf node for z with an internal node having x and y as children, represents an optimal prefix code for the alphabet C .

Theorem 16.4

Procedure HUFFMAN produces an optimal prefix code.

1.5 Overview of the Formalization

This report presents a formalization of the proof of Huffman’s algorithm written using Isabelle/HOL [9]. Our proof is based on the informal proofs given by Knuth and Cormen et al. The development was done independently of Laurent Théry’s Coq proof [11, 12], which through its “cover” concept represents a considerable departure from the standard proof.

The development consists of 90 lemmas and 5 theorems. Most of them have very short proofs thanks to the extensive use of simplification rules and custom induction rules. The remaining proofs are written using the structured proof format Isar [13] and are accompanied by informal arguments and diagrams.

The report is organized as follows. Section 2 defines the datatypes for binary code trees and forests and develops a small library of related functions. (Incidentally, there is nothing special about binary codes and binary trees. Huffman’s

algorithm and its proof can be generalized to n -ary trees [6, p. 405 and 595].) Section 3 presents a functional implementation of the algorithm. Section 4 defines several tree manipulation functions needed for the proof. Section 5 presents three key lemmas and concludes with the optimality theorem. Section 6 compares our work with Théry’s Coq proof. Finally, Section 7 concludes the report.

1.6 Overview of Isabelle’s HOL Logic

This section presents a brief overview of the Isabelle/HOL logic, so that readers not familiar with the system can at least understand the lemmas and theorems, if not the proofs. Readers who already know Isabelle are encouraged to skip this section.

Isabelle is a generic theorem prover whose built-in metalogic is a fragment of higher-order logic [4, 9]. The metalogical operators are material implication, written $[\![\varphi_1; \dots; \varphi_n]\!] \implies \psi$ (“if φ_1 and \dots and φ_n , then ψ ”), universal quantification, written $\bigwedge x_1 \dots x_n. \psi$ (“for all x_1, \dots, x_n we have ψ ”), and equality, written $t \equiv u$.

The incarnation of Isabelle that we use in this development, Isabelle/HOL, provides a more elaborate version of higher-order logic, complete with the familiar connectives and quantifiers (\neg , \wedge , \vee , \longrightarrow , \forall , and \exists) on terms of type *bool*. In addition, $=$ expresses equivalence. The formulas $\bigwedge x_1 \dots x_m. [\![\varphi_1; \dots; \varphi_n]\!] \implies \psi$ and $\forall x_1. \dots \forall x_m. \varphi_1 \wedge \dots \wedge \varphi_n \longrightarrow \psi$ are logically equivalent, but they interact differently with Isabelle’s proof tactics.

The term language consists of simply typed λ -terms written in an ML-like syntax [8]. In particular, function application expects no parentheses around the argument list and no commas between the arguments, as in $f x y$. Syntactic sugar provides an infix syntax for common operators, such as $x = y$ and $x + y$. Types are inferred automatically in most cases, but they can always be supplied using an annotation $t::\tau$, where t is a term and τ is its type. The type of total functions from α to β is written $\alpha \Rightarrow \beta$. Variables may range over functions.

The type of natural numbers is called *nat*. The type of lists over type α , written α *list*, features the empty list $[\]$, the infix constructor $x \cdot xs$ (where x is an element of type α and xs is a list over α), and the conversion function *set* from lists to sets. The type of sets over α is written α *set*. Operations on sets are written using traditional mathematical notations.

1.7 Head of the Theory File

The Isabelle theory starts in the standard way.

```
theory Huffman
imports Main
begin
```

We attach the *simp* attribute to some predefined lemmas to add them to the de-

fault set of simplification rules.

```
declare Int_Un_distrib [simp]  
         Int_Un_distrib2 [simp]  
         min_max.sup_absorb1 [simp]  
         min_max.sup_absorb2 [simp]
```

2 Definition of Prefix Code Trees and Forests

2.1 Tree Datatype

A *prefix code tree* is a full binary tree in which leaf nodes are of the form *Leaf w a*, where *a* is a symbol and *w* is the frequency associated with *a*, and inner nodes are of the form *InnerNode w t₁ t₂*, where *t₁* and *t₂* are the left and right subtrees and *w* caches the sum of the weights of *t₁* and *t₂*. Prefix code trees are polymorphic on the symbol datatype α .

```
datatype  $\alpha$  tree =  
  Leaf nat  $\alpha$   
  InnerNode nat ( $\alpha$  tree) ( $\alpha$  tree)
```

2.2 Forest Datatype

The intermediate steps of Huffman's algorithm involve a list of prefix code trees, or *prefix code forest*.

```
types  $\alpha$  forest =  $\alpha$  tree list
```

2.3 Alphabet

The *alphabet* of a code tree is the set of symbols appearing in the tree's leaf nodes.

```
primrec alphabet ::  $\alpha$  tree  $\Rightarrow$   $\alpha$  set where  
alphabet (Leaf w a) = {a}  
alphabet (InnerNode w t1 t2) = alphabet t1  $\cup$  alphabet t2
```

For set and predicates, Isabelle gives us the choice between inductive definitions (**inductive_set** and **inductive**) and recursive functions (**primrec**, **fun**, and **function**). In this development, we consistently favor recursion over induction, for two reasons:

- Recursion gives rise to simplification rules that greatly help automatic proof tactics. In contrast, reasoning about inductively defined sets and predicates involves introduction and elimination rules, which are more clumsy than simplification rules.

- Isabelle’s counterexample generator **quickcheck** [2], which is very useful when developing proofs top down (together with **sorry**), has limited support for inductive definitions.

The alphabet of a forest is defined as the union of the alphabets of the trees that compose it. Although Isabelle supports overloading for non-overlapping types, we avoid many type inference problems by attaching an ‘ f ’ subscript to the forest generalizations of functions defined on trees.

primrec $alphabet_f :: \alpha \text{ forest} \Rightarrow \alpha \text{ set}$ **where**
 $alphabet_f [] = \emptyset$
 $alphabet_f (t \cdot ts) = alphabet\ t \cup alphabet_f\ ts$

Alphabets are central to our proofs, and we need the following basic facts about them.

lemma $finite_alphabet$ [simp]:
 $finite\ (alphabet\ t)$
 $\langle proof \rangle$

lemma $exists_in_alphabet$:
 $\exists a. a \in alphabet\ t$
 $\langle proof \rangle$

2.4 Consistency

A tree is *consistent* if for each inner node the alphabet of the two subtrees are disjoint. Intuitively, this means that every symbol in the alphabet occurs in exactly one leaf node. Although this well-formedness property isn’t mentioned in algorithms textbooks [1, 3, 6], it is essential and appears as an assumption in many of our lemmas.

primrec $consistent :: \alpha \text{ tree} \Rightarrow bool$ **where**
 $consistent\ (Leaf\ w\ a) = True$
 $consistent\ (InnerNode\ w\ t_1\ t_2) =$
 $(consistent\ t_1 \wedge consistent\ t_2 \wedge alphabet\ t_1 \cap alphabet\ t_2 = \emptyset)$

primrec $consistent_f :: \alpha \text{ forest} \Rightarrow bool$ **where**
 $consistent_f\ [] = True$
 $consistent_f\ (t \cdot ts) =$
 $(consistent\ t \wedge consistent_f\ ts \wedge alphabet\ t \cap alphabet_f\ ts = \emptyset)$

Several of our proofs are by structural induction on consistent trees t and involve one symbol a . These proofs typically distinguish the following cases.

CONSISTENCY CHECK: The property holds if t is inconsistent. (This is always the case when $consistent\ t$ occurs among the property’s assumptions.)

BASE CASE: $t = \text{Leaf } w \ b$.

INDUCTION STEP: $t = \text{InnerNode } w \ t_1 \ t_2$.

SUBCASE 1: a belongs to t_1 but not to t_2 .

SUBCASE 2: a belongs to t_2 but not to t_1 .

SUBCASE 3: a belongs to neither t_1 nor t_2 .

Thanks to the consistency check, we can rule out the subcase where a belongs to both subtrees.

Instead of performing the above case distinction manually, we encode it in a custom induction rule. This saves us from writing repetitive proof scripts and helps Isabelle's automatic proof tactics.

lemma *tree_induct_consistent*:

$$\begin{aligned} & \llbracket \bigwedge t \ a. \ \neg \text{consistent } t \implies P \ t \ a; \\ & \quad \bigwedge w_b \ b \ a. \ P \ (\text{Leaf } w_b \ b) \ a; \\ & \quad \bigwedge w \ t_1 \ t_2 \ a. \\ & \quad \quad \llbracket \text{consistent } t_1; \text{consistent } t_2; \text{alphabet } t_1 \cap \text{alphabet } t_2 = \emptyset; \\ & \quad \quad \quad a \in \text{alphabet } t_1; a \notin \text{alphabet } t_2; P \ t_1 \ a; P \ t_2 \ a \rrbracket \implies \\ & \quad \quad P \ (\text{InnerNode } w \ t_1 \ t_2) \ a; \\ & \quad \bigwedge w \ t_1 \ t_2 \ a. \\ & \quad \quad \llbracket \text{consistent } t_1; \text{consistent } t_2; \text{alphabet } t_1 \cap \text{alphabet } t_2 = \emptyset; \\ & \quad \quad \quad a \notin \text{alphabet } t_1; a \in \text{alphabet } t_2; P \ t_1 \ a; P \ t_2 \ a \rrbracket \implies \\ & \quad \quad P \ (\text{InnerNode } w \ t_1 \ t_2) \ a; \\ & \quad \bigwedge w \ t_1 \ t_2 \ a. \\ & \quad \quad \llbracket \text{consistent } t_1; \text{consistent } t_2; \text{alphabet } t_1 \cap \text{alphabet } t_2 = \emptyset; \\ & \quad \quad \quad a \notin \text{alphabet } t_1; a \notin \text{alphabet } t_2; P \ t_1 \ a; P \ t_2 \ a \rrbracket \implies \\ & \quad \quad P \ (\text{InnerNode } w \ t_1 \ t_2) \ a \rrbracket \implies \\ & P \ t \ a \langle \text{proof} \rangle \end{aligned}$$

2.5 Symbol Depths

The *depth* of a symbol is the length of the path from the root to the leaf node labeled with that symbol. Symbols that don't occur in the tree or that occur at the root of a one-node tree have depth 0. If a symbol occurs in several leaf nodes (which may happen with inconsistent trees), the depth is arbitrarily defined in terms of the leftmost node labeled with that symbol.

primrec *depth* :: $\alpha \ \text{tree} \Rightarrow \alpha \Rightarrow \text{nat}$ **where**

depth (*Leaf* $w \ b$) $a = 0$

depth (*InnerNode* $w \ t_1 \ t_2$) $a =$

(if $a \in \text{alphabet } t_1$ then *depth* $t_1 \ a + 1$

else if $a \in \text{alphabet } t_2$ then *depth* $t_2 \ a + 1$

else 0)

2.6 Height

The *height* of a tree is the length of the longest path from the root to a leaf node. This is readily generalized to forests by taking the maximum of the trees' heights. Note that a tree has height 0 if and only if it is a leaf node, and that a forest has height 0 if and only if all its trees are leaf nodes.

primrec *height* :: α tree \Rightarrow nat **where**
height (Leaf *w a*) = 0
height (InnerNode *w t₁ t₂*) = max (*height t₁*) (*height t₂*) + 1

primrec *height_F* :: α forest \Rightarrow nat **where**
height_F [] = 0
height_F (*t · ts*) = max (*height t*) (*height_F ts*)

The depth of any symbol in the tree is bounded by the tree's height, and there exists a symbol with a depth equal to the height.

lemma *depth_le_height*:
depth t a \leq *height t*
 \langle proof \rangle

lemma *exists_at_height*:
consistent t $\implies \exists a \in$ alphabet *t*. *depth t a* = *height t*
 \langle proof \rangle

The following elimination rules help Isabelle's classical prover, notably the **auto** tactic. They are easy consequences of the inequation *depth t a* \leq *height t*.

lemma *depth_max_heightE_left* [elim!]:
 \llbracket *depth t₁ a* = max (*height t₁*) (*height t₂*);
 \llbracket *depth t₁ a* = *height t₁*; *height t₁* \geq *height t₂* $\rrbracket \implies P \implies$
P
 \langle proof \rangle

lemma *depth_max_heightE_right* [elim!]:
 \llbracket *depth t₂ a* = max (*height t₁*) (*height t₂*);
 \llbracket *depth t₂ a* = *height t₂*; *height t₂* \geq *height t₁* $\rrbracket \implies P \implies$
P
 \langle proof \rangle

We also need the following lemma.

lemma *height_gt_0_alphabet_eq_imp_height_gt_0*:
assumes *height t* > 0 *consistent t* *alphabet t* = *alphabet u*
shows *height u* > 0
 \langle proof \rangle

2.7 Symbol Frequencies

The *frequency* of a symbol is the sum of the weights attached to the leaf nodes labeled with that symbol. If the tree is consistent, the sum comprises at most one nonzero term. The generalization to forests is straightforward. If two trees have the same alphabet and symbol frequencies, we say that they are *compatible*.

primrec $freq :: \alpha \text{ tree} \Rightarrow \alpha \Rightarrow \text{nat}$ **where**
 $freq (\text{Leaf } w \ a) = (\lambda b. \text{if } b = a \text{ then } w \text{ else } 0)$
 $freq (\text{InnerNode } w \ t_1 \ t_2) = (\lambda b. freq \ t_1 \ b + freq \ t_2 \ b)$

primrec $freq_F :: \alpha \text{ forest} \Rightarrow \alpha \Rightarrow \text{nat}$ **where**
 $freq_F [] = (\lambda b. 0)$
 $freq_F (t \cdot ts) = (\lambda b. freq \ t \ b + freq_F \ ts \ b)$

Alphabet and symbol frequencies are intimately related. Simplification rules ensure that sums of the form $freq \ t_1 \ a + freq \ t_2 \ a$ collapse to a single term when we know which tree a belongs to.

lemma $notin_alphabet_imp_freq_0$ [simp]:
 $a \notin \text{alphabet } t \implies freq \ t \ a = 0$
 <proof>

lemma $notin_alphabet_F_imp_freq_F_0$ [simp]:
 $a \notin \text{alphabet}_F \ ts \implies freq_F \ ts \ a = 0$
 <proof>

lemma $freq_0_right$ [simp]:
 $\llbracket \text{alphabet } t_1 \cap \text{alphabet } t_2 = \emptyset; a \in \text{alphabet } t_1 \rrbracket \implies freq \ t_2 \ a = 0$
 <proof>

lemma $freq_0_left$ [simp]:
 $\llbracket \text{alphabet } t_1 \cap \text{alphabet } t_2 = \emptyset; a \in \text{alphabet } t_2 \rrbracket \implies freq \ t_1 \ a = 0$
 <proof>

We close this section with a more technical lemma.

lemma $height_F_0_imp_Leaf_freq_in_set$:
 $\llbracket \text{consistent}_F \ ts; \text{height}_F \ ts = 0; a \in \text{alphabet}_F \ ts \rrbracket \implies$
 $\text{Leaf } (freq_F \ ts \ a) \ a \in \text{set } ts$
 <proof>

2.8 Weight

The *weight* function returns the weight of a tree. In the *InnerNode* case, we ignore the weight cached in the node and instead compute the tree's weight recursively. This is more robust than relying on the cache and simplifies reasoning.

primrec $weight :: \alpha\ tree \Rightarrow nat$ **where**

$weight\ (Leaf\ w\ a) = w$

$weight\ (InnerNode\ w\ t_1\ t_2) = weight\ t_1 + weight\ t_2$

The weight of a tree is the sum of the frequencies of its symbols.

lemma $weight_eq_Sum_freq$:

$consistent\ t \Longrightarrow weight\ t = \sum_{a \in alphabet\ t} freq\ t\ a$

<proof>

The assumption $consistent\ t$ is not necessary, but it simplifies the proof by letting us invoke the lemma $setsum_Un_disjoint$:

$$\llbracket finite\ A; finite\ B; A \cap B = \emptyset \rrbracket \Longrightarrow \sum_{x \in A} g\ x + \sum_{x \in B} g\ x = \sum_{x \in A \cup B} g\ x.$$

2.9 Cost

The cost of a consistent tree (sometimes called the *weighted path length*) is given by the sum $\sum_{a \in alphabet\ t} freq\ t\ a \times depth\ t\ a$. It obeys a simple recursive law.

primrec $cost :: \alpha\ tree \Rightarrow nat$ **where**

$cost\ (Leaf\ w\ a) = 0$

$cost\ (InnerNode\ w\ t_1\ t_2) = weight\ t_1 + cost\ t_1 + weight\ t_2 + cost\ t_2$

One interpretation of this recursive law is that the cost of a tree is the sum of the weights of its inner nodes [6, p. 405]. (Recall that $weight\ (InnerNode\ w\ t_1\ t_2) = weight\ t_1 + weight\ t_2$.) Since the cost of a tree is such a fundamental concept, it seems necessary to prove that the above function definition is correct.

theorem $cost_eq_Sum_freq_mult_depth$:

$consistent\ t \Longrightarrow cost\ t = \sum_{a \in alphabet\ t} freq\ t\ a \times depth\ t\ a$

The proof is by structural induction on t . If $t = Leaf\ w\ b$, both sides of the equation simplify to 0. This leaves the case $t = InnerNode\ w\ t_1\ t_2$. Let A , A_1 , and A_2 stand

for *alphabet* t , *alphabet* t_1 , and *alphabet* t_2 , respectively. We have

$$\begin{aligned}
& \text{cost } t \\
= & \hspace{15em} \text{(definition of } \text{cost}) \\
& \text{weight } t_1 + \text{cost } t_1 + \text{weight } t_2 + \text{cost } t_2 \\
= & \hspace{15em} \text{(induction hypothesis)} \\
& \text{weight } t_1 + \sum_{a \in A_1} \text{freq } t_1 a \times \text{depth } t_1 a + \\
& \text{weight } t_2 + \sum_{a \in A_2} \text{freq } t_2 a \times \text{depth } t_2 a \\
= & \hspace{15em} \text{(definition of } \text{depth}, \text{ consistency)} \\
& \text{weight } t_1 + \sum_{a \in A_1} \text{freq } t_1 a \times (\text{depth } t a - 1) + \\
& \text{weight } t_2 + \sum_{a \in A_2} \text{freq } t_2 a \times (\text{depth } t a - 1) \\
= & \hspace{15em} \text{(distributivity of } \times \text{ and } \Sigma \text{ over } -) \\
& \text{weight } t_1 + \sum_{a \in A_1} \text{freq } t_1 a \times \text{depth } t a - \sum_{a \in A_1} \text{freq } t_1 a + \\
& \text{weight } t_2 + \sum_{a \in A_2} \text{freq } t_2 a \times \text{depth } t a - \sum_{a \in A_2} \text{freq } t_2 a \\
= & \hspace{15em} \text{(weight_eq_Sum_freq)} \\
& \sum_{a \in A_1} \text{freq } t_1 a \times \text{depth } t a + \sum_{a \in A_2} \text{freq } t_2 a \times \text{depth } t a \\
= & \hspace{15em} \text{(definition of } \text{freq}, \text{ consistency)} \\
& \sum_{a \in A_1} \text{freq } t a \times \text{depth } t a + \sum_{a \in A_2} \text{freq } t a \times \text{depth } t a \\
= & \hspace{15em} \text{(setsum_Un_disjoint, consistency)} \\
& \sum_{a \in A_1 \cup A_2} \text{freq } t a \times \text{depth } t a \\
= & \hspace{15em} \text{(definition of } \text{alphabet}) \\
& \sum_{a \in A} \text{freq } t a \times \text{depth } t a.
\end{aligned}$$

The structured proof closely follows this argument.

<proof>

Finally, it should come as no surprise that trees with height 0 have cost 0.

lemma *height_0_imp_cost_0* [simp]:

$\text{height } t = 0 \implies \text{cost } t = 0$

<proof>

2.10 Optimality

A tree is optimum if and only if its cost is not greater than the cost of any compatible tree. We can ignore inconsistent trees without loss of generality.

definition *optimum* :: $\alpha \text{ tree} \Rightarrow \text{bool}$ **where**

$\text{optimum } t \equiv$

$\forall u. \text{consistent } u \longrightarrow \text{alphabet } t = \text{alphabet } u \longrightarrow \text{freq } t = \text{freq } u \longrightarrow$
 $\text{cost } t \leq \text{cost } u$

3 Functional Implementation of Huffman's Algorithm

3.1 Cached Weight

The *cached weight* of a node is the weight stored directly in the node. Our arguments rely on the computed weight (embodied by the *weight* function) rather than the cached weight, but the implementation of Huffman's algorithm uses the cached weight for performance reasons.

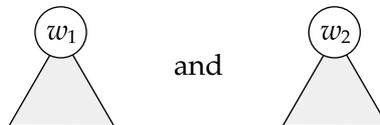
primrec *cachedWeight* :: α tree \Rightarrow nat **where**
cachedWeight (Leaf w a) = w
cachedWeight (InnerNode w t_1 t_2) = w

The cached weight of a leaf node is identical to its computed weight.

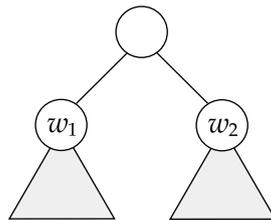
lemma *height_0_imp_cachedWeight_eq_weight* [simp]:
height t = 0 \implies *cachedWeight* t = *weight* t
 <proof>

3.2 Tree Union

The implementation of Huffman's algorithm builds on two additional auxiliary functions. The first one, *uniteTrees*, takes two trees



and returns the tree



definition *uniteTrees* :: α tree \Rightarrow α tree \Rightarrow α tree **where**
uniteTrees t_1 t_2 \equiv InnerNode (*cachedWeight* t_1 + *cachedWeight* t_2) t_1 t_2

The alphabet, consistency, and symbol frequencies of a united tree are easy to connect to the homologous properties of the subtrees.

lemma *alphabet_uniteTrees* [simp]:
alphabet (*uniteTrees* t_1 t_2) = *alphabet* t_1 \cup *alphabet* t_2
 <proof>

lemma *consistent_uniteTrees* [simp]:
 $\llbracket \text{consistent } t_1; \text{consistent } t_2; \text{alphabet } t_1 \cap \text{alphabet } t_2 = \emptyset \rrbracket \implies$
 $\text{consistent } (\text{uniteTrees } t_1 \ t_2)$
 ⟨proof⟩

lemma *freq_uniteTrees* [simp]:
 $\text{freq } (\text{uniteTrees } t_1 \ t_2) = (\lambda a. \text{freq } t_1 \ a + \text{freq } t_2 \ a)$
 ⟨proof⟩

3.3 Ordered Tree Insertion

The auxiliary function *insortTree* inserts a tree into a forest sorted by cached weight, preserving the sort order.

primrec *insortTree* :: $\alpha \text{ tree} \Rightarrow \alpha \text{ forest} \Rightarrow \alpha \text{ forest}$ **where**
 $\text{insortTree } u \ [] = [u]$
 $\text{insortTree } u \ (t \cdot ts) =$
 (if $\text{cachedWeight } u \leq \text{cachedWeight } t$ then $u \cdot t \cdot ts$
 else $t \cdot \text{insortTree } u \ ts$)

The resulting forest contains one more tree than the original forest. Clearly, it cannot be empty.

lemma *length_insortTree* [simp]:
 $\text{length } (\text{insortTree } t \ ts) = \text{length } ts + 1$
 ⟨proof⟩

lemma *insortTree_ne_Nil* [simp]:
 $\text{insortTree } t \ ts \neq []$
 ⟨proof⟩

The alphabet, consistency, symbol frequencies, and height of a forest after insertion are easy to relate to the homologous properties of the original forest and the inserted tree.

lemma *alphabet_F_insortTree* [simp]:
 $\text{alphabet}_F (\text{insortTree } t \ ts) = \text{alphabet } t \cup \text{alphabet}_F \ ts$
 ⟨proof⟩

lemma *consistent_F_insortTree* [simp]:
 $\text{consistent}_F (\text{insortTree } t \ ts) = \text{consistent}_F (t \cdot ts)$
 ⟨proof⟩

lemma *freq_F_insortTree* [simp]:
 $\text{freq}_F (\text{insortTree } t \ ts) = (\lambda a. \text{freq } t \ a + \text{freq}_F \ ts \ a)$
 ⟨proof⟩

lemma *height_F_insortTree* [simp]:
 $\text{height}_F (\text{insortTree } t \ ts) = \max (\text{height } t) (\text{height}_F \ ts)$
 ⟨proof⟩

3.4 The Main Algorithm

Huffman’s algorithm repeatedly unites the first two trees of the forest it receives as argument until a single tree is left. It should initially be invoked with a list of leaf nodes sorted by weight.

fun *huffman* :: α forest \Rightarrow α tree **where**
huffman [t] = t
huffman (t₁ · t₂ · ts) = *huffman* (insortTree (uniteTrees t₁ t₂) ts)

The tree returned by the algorithm preserves the alphabet, consistency, and symbol frequencies of the original forest.

theorem *alphabet_huffman* [simp]:
 $ts \neq [] \implies \text{alphabet} (\text{huffman } ts) = \text{alphabet}_F \ ts$
 ⟨proof⟩

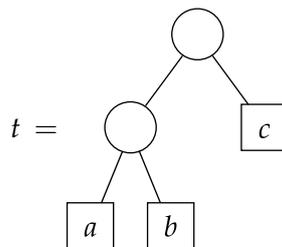
theorem *consistent_huffman* [simp]:
 $\llbracket \text{consistent}_F \ ts; \ ts \neq [] \rrbracket \implies \text{consistent} (\text{huffman } ts)$
 ⟨proof⟩

theorem *freq_huffman* [simp]:
 $ts \neq [] \implies \text{freq} (\text{huffman } ts) = \text{freq}_F \ ts$
 ⟨proof⟩

4 Definition of Auxiliary Functions Used in the Proof

4.1 Sibling of a Symbol

The *sibling* of a symbol *a* in a tree *t* is the label of the node that is the (left or right) sibling of the node labeled with *a* in *t*. If the symbol *a* is not in *t*’s alphabet or it occurs in a node with no sibling leaf, we define the sibling as being *a* itself. Thus, we have *sibling* *t* *a* = *b*, *sibling* *t* *b* = *a*, and *sibling* *t* *c* = *c* for the tree



```

fun sibling ::  $\alpha$  tree  $\Rightarrow$   $\alpha$   $\Rightarrow$   $\alpha$  where
sibling (Leaf  $w_b$   $b$ )  $a$  =  $a$ 
sibling (InnerNode  $w$  (Leaf  $w_b$   $b$ ) (Leaf  $w_c$   $c$ ))  $a$  =
  (if  $a = b$  then  $c$  else if  $a = c$  then  $b$  else  $a$ )
sibling (InnerNode  $w$   $t_1$   $t_2$ )  $a$  =
  (if  $a \in \text{alphabet } t_1$  then sibling  $t_1$   $a$ 
   else if  $a \in \text{alphabet } t_2$  then sibling  $t_2$   $a$ 
   else  $a$ )

```

Because *sibling* is defined using sequential pattern matching [7], reasoning about it can become tedious. Simplification rules therefore play an important role.

```

lemma notin_alphabet_imp_sibling_id [simp]:
 $a \notin \text{alphabet } t \implies \text{sibling } t \ a = a$ 
<proof>

```

```

lemma height_0_imp_sibling_id [simp]:
height  $t = 0 \implies \text{sibling } t \ a = a$ 
<proof>

```

```

lemma height_gt_0_in_alphabet_imp_sibling_left [simp]:
[[height  $t_1 > 0$ ;  $a \in \text{alphabet } t_1$ ]]  $\implies$ 
sibling (InnerNode  $w$   $t_1$   $t_2$ )  $a = \text{sibling } t_1 \ a$ 
<proof>

```

```

lemma height_gt_0_in_alphabet_imp_sibling_right [simp]:
[[height  $t_2 > 0$ ;  $a \in \text{alphabet } t_1$ ]]  $\implies$ 
sibling (InnerNode  $w$   $t_1$   $t_2$ )  $a = \text{sibling } t_1 \ a$ 
<proof>

```

```

lemma height_gt_0_notin_alphabet_imp_sibling_left [simp]:
[[height  $t_1 > 0$ ;  $a \notin \text{alphabet } t_1$ ]]  $\implies$ 
sibling (InnerNode  $w$   $t_1$   $t_2$ )  $a = \text{sibling } t_2 \ a$ 
<proof>

```

```

lemma height_gt_0_notin_alphabet_imp_sibling_right [simp]:
[[height  $t_2 > 0$ ;  $a \notin \text{alphabet } t_1$ ]]  $\implies$ 
sibling (InnerNode  $w$   $t_1$   $t_2$ )  $a = \text{sibling } t_2 \ a$ 
<proof>

```

```

lemma either_height_gt_0_imp_sibling [simp]:
height  $t_1 > 0 \vee \text{height } t_2 > 0 \implies$ 
sibling (InnerNode  $w$   $t_1$   $t_2$ )  $a =$ 
  (if  $a \in \text{alphabet } t_1$  then sibling  $t_1$   $a$  else sibling  $t_2$   $a$ )
<proof>

```

The following rules are also useful for reasoning about siblings and alphabets.

lemma *in_alphabet_imp_sibling_in_alphabet*:

$a \in \text{alphabet } t \implies \text{sibling } t a \in \text{alphabet } t$

<proof>

lemma *sibling_ne_imp_sibling_in_alphabet*:

$\text{sibling } t a \neq a \implies \text{sibling } t a \in \text{alphabet } t$

<proof>

The default induction rule for *sibling* distinguishes four cases.

BASE CASE: $t = \text{Leaf } w b$.

INDUCTION STEP 1: $t = \text{InnerNode } w (\text{Leaf } w_b b) (\text{Leaf } w_c c)$.

INDUCTION STEP 2: $t = \text{InnerNode } w (\text{InnerNode } w_1 t_{11} t_{12}) t_2$.

INDUCTION STEP 3: $t = \text{InnerNode } w t_1 (\text{InnerNode } w_2 t_{21} t_{22})$.

This rule leaves much to be desired. First, the last two cases overlap and can normally be handled the same way, so they should be combined. Second, the nested *InnerNode* constructors in the last two cases reduce readability. Third, under the assumption that t is consistent, we would like to perform the same case distinction on a as we did for *tree_induct_consistent*, with the same benefits for automation.

These observations lead us to develop a custom induction rule that distinguishes the following cases.

CONSISTENCY CHECK: The property holds if t is inconsistent.

BASE CASE: $t = \text{Leaf } w b$.

INDUCTION STEP 1: $t = \text{InnerNode } w (\text{Leaf } w_b b) (\text{Leaf } w_c c)$ with $b \neq c$.

INDUCTION STEP 2: $t = \text{InnerNode } w t_1 t_2$ and either t_1 or t_2 has nonzero height.

SUBCASE 1: a belongs to t_1 but not to t_2 .

SUBCASE 2: a belongs to t_2 but not to t_1 .

SUBCASE 3: a belongs to neither t_1 nor t_2 .

The statement of the rule and its proof are similar to what we did for consistent trees, the main difference being that we now have two induction steps instead of one.

lemma *sibling_induct_consistent*:

$\llbracket \bigwedge t a. \neg \text{consistent } t \implies P t a;$

$\bigwedge w b a. P (\text{Leaf } w b) a;$

$$\begin{aligned}
& \wedge w w_b b w_c c a. b \neq c \implies P (\text{InnerNode } w (\text{Leaf } w_b b) (\text{Leaf } w_c c)) a; \\
& \wedge w t_1 t_2 a. \\
& \quad \llbracket \text{consistent } t_1; \text{ consistent } t_2; \text{ alphabet } t_1 \cap \text{alphabet } t_2 = \emptyset; \\
& \quad \quad \text{height } t_1 > 0 \vee \text{height } t_2 > 0; a \in \text{alphabet } t_1; \\
& \quad \quad \text{sibling } t_1 a \in \text{alphabet } t_1; a \notin \text{alphabet } t_2; \\
& \quad \quad \text{sibling } t_1 a \notin \text{alphabet } t_2; P t_1 a \rrbracket \implies \\
& \quad P (\text{InnerNode } w t_1 t_2) a; \\
& \wedge w t_1 t_2 a. \\
& \quad \llbracket \text{consistent } t_1; \text{ consistent } t_2; \text{ alphabet } t_1 \cap \text{alphabet } t_2 = \emptyset; \\
& \quad \quad \text{height } t_1 > 0 \vee \text{height } t_2 > 0; a \notin \text{alphabet } t_1; \\
& \quad \quad \text{sibling } t_2 a \notin \text{alphabet } t_1; a \in \text{alphabet } t_2; \\
& \quad \quad \text{sibling } t_2 a \in \text{alphabet } t_2; P t_2 a \rrbracket \implies \\
& \quad P (\text{InnerNode } w t_1 t_2) a; \\
& \wedge w t_1 t_2 a. \\
& \quad \llbracket \text{consistent } t_1; \text{ consistent } t_2; \text{ alphabet } t_1 \cap \text{alphabet } t_2 = \emptyset; \\
& \quad \quad \text{height } t_1 > 0 \vee \text{height } t_2 > 0; a \notin \text{alphabet } t_1; a \notin \text{alphabet } t_2 \rrbracket \implies \\
& \quad P (\text{InnerNode } w t_1 t_2) a \rrbracket \implies \\
& P t a \\
& \langle \text{proof} \rangle
\end{aligned}$$

The custom induction rule allows us to prove new properties of *sibling* with little effort.

lemma *sibling_sibling_id* [simp]:
 $\text{consistent } t \implies \text{sibling } t (\text{sibling } t a) = a$
 $\langle \text{proof} \rangle$

lemma *sibling_reciprocal*:
 $\llbracket \text{consistent } t; \text{sibling } t a = b \rrbracket \implies \text{sibling } t b = a$
 $\langle \text{proof} \rangle$

lemma *depth_height_imp_sibling_ne*:
 $\llbracket \text{depth } t a = \text{height } t; \text{ consistent } t; \text{height } t > 0; a \in \text{alphabet } t \rrbracket \implies$
 $\text{sibling } t a \neq a$
 $\langle \text{proof} \rangle$

lemma *depth_sibling* [simp]:
 $\text{consistent } t \implies \text{depth } t (\text{sibling } t a) = \text{depth } t a$
 $\langle \text{proof} \rangle$

4.2 Leaf Interchange

The *swapLeaves* function takes a tree t together with two symbols a, b and their frequencies w_a, w_b , and returns the tree t in which the leaf nodes labeled with a and b are exchanged. When invoking *swapLeaves*, we normally pass *freq* t a and

$\text{freq } t \ b$ for w_a and w_b .

Notice that we do not bother updating the cached weight of the ancestor nodes when performing the interchange. The cached weight is used only in the implementation of Huffman's algorithm, which doesn't invoke swapLeaves .

primrec $\text{swapLeaves} :: \alpha \ \text{tree} \Rightarrow \text{nat} \Rightarrow \alpha \Rightarrow \text{nat} \Rightarrow \alpha \Rightarrow \alpha \ \text{tree}$ **where**
 $\text{swapLeaves} (\text{Leaf } w_c \ c) \ w_a \ a \ w_b \ b =$
 (if $c = a$ then $\text{Leaf } w_b \ b$ else if $c = b$ then $\text{Leaf } w_a \ a$ else $\text{Leaf } w_c \ c$)
 $\text{swapLeaves} (\text{InnerNode } w \ t_1 \ t_2) \ w_a \ a \ w_b \ b =$
 $\text{InnerNode } w \ (\text{swapLeaves } t_1 \ w_a \ a \ w_b \ b) \ (\text{swapLeaves } t_2 \ w_a \ a \ w_b \ b)$

Swapping a symbol a with itself leaves the tree t unchanged if a does not belong to it or if the specified frequencies w_a and w_b equal $\text{freq } t \ a$.

lemma $\text{swapLeaves_id_when_notin_alphabet}$ [simp]:
 $a \notin \text{alphabet } t \Longrightarrow \text{swapLeaves } t \ w \ a \ w' \ a = t$
 <proof>

lemma swapLeaves_id [simp]:
 $\text{consistent } t \Longrightarrow \text{swapLeaves } t \ (\text{freq } t \ a) \ a \ (\text{freq } t \ a) \ a = t$
 <proof>

The alphabet, consistency, symbol depths, height, and symbol frequencies of the tree $\text{swapLeaves } t \ w_a \ a \ w_b \ b$ can be related to the homologous properties of t .

lemma $\text{alphabet_swapLeaves}$:
 $\text{alphabet} (\text{swapLeaves } t \ w_a \ a \ w_b \ b) =$
 (if $a \in \text{alphabet } t$ then
 if $b \in \text{alphabet } t$ then $\text{alphabet } t$ else $(\text{alphabet } t - \{a\}) \cup \{b\}$
 else
 if $b \in \text{alphabet } t$ then $(\text{alphabet } t - \{b\}) \cup \{a\}$ else $\text{alphabet } t$)
 <proof>

lemma $\text{consistent_swapLeaves}$ [simp]:
 $\text{consistent } t \Longrightarrow \text{consistent} (\text{swapLeaves } t \ w_a \ a \ w_b \ b)$
 <proof>

lemma $\text{depth_swapLeaves_neither}$ [simp]:
 $\llbracket \text{consistent } t; c \neq a; c \neq b \rrbracket \Longrightarrow \text{depth} (\text{swapLeaves } t \ w_a \ a \ w_b \ b) \ c = \text{depth } t \ c$
 <proof>

lemma height_swapLeaves [simp]:
 $\text{height} (\text{swapLeaves } t \ w_a \ a \ w_b \ b) = \text{height } t$
 <proof>

lemma freq_swapLeaves [simp]:
 $\llbracket \text{consistent } t; a \neq b \rrbracket \Longrightarrow$

$$\begin{aligned} \text{freq } (\text{swapLeaves } t \ w_a \ a \ w_b \ b) = & \\ & (\lambda c. \text{ if } c = a \text{ then if } b \in \text{alphabet } t \text{ then } w_a \text{ else } 0 \\ & \quad \text{else if } c = b \text{ then if } a \in \text{alphabet } t \text{ then } w_b \text{ else } 0 \\ & \quad \text{else freq } t \ c) \\ \langle \text{proof} \rangle & \end{aligned}$$

For the lemmas concerned with the resulting tree's weight and cost, we avoid subtraction on natural numbers by rearranging terms. For example, we write

$$\text{weight } (\text{swapLeaves } t \ w_a \ a \ w_b \ b) + \text{freq } t \ a = \text{weight } t + w_b$$

rather than the more conventional

$$\text{weight } (\text{swapLeaves } t \ w_a \ a \ w_b \ b) = \text{weight } t + w_b - \text{freq } t \ a.$$

In Isabelle/HOL, these two equations are not equivalent, because by definition $m - n = 0$ if $n > m$. We could use the second equation and additionally assert that $\text{freq } t \ a \leq \text{weight } t$ (an easy consequence of *weight_eq_Sum_freq*), and then apply the **arith** tactic, but it is much simpler to use the first equation and stay with **simp** and **auto**. Another option would be to use integers instead of natural numbers.

lemma *weight_swapLeaves*:

$$\begin{aligned} \llbracket \text{consistent } t; a \neq b \rrbracket \implies & \\ \text{if } a \in \text{alphabet } t \text{ then} & \\ \quad \text{if } b \in \text{alphabet } t \text{ then} & \\ \quad \quad \text{weight } (\text{swapLeaves } t \ w_a \ a \ w_b \ b) + \text{freq } t \ a + \text{freq } t \ b = & \\ \quad \quad \quad \text{weight } t + w_a + w_b & \\ \quad \text{else} & \\ \quad \quad \text{weight } (\text{swapLeaves } t \ w_a \ a \ w_b \ b) + \text{freq } t \ a = \text{weight } t + w_b & \\ \text{else} & \\ \quad \text{if } b \in \text{alphabet } t \text{ then} & \\ \quad \quad \text{weight } (\text{swapLeaves } t \ w_a \ a \ w_b \ b) + \text{freq } t \ b = \text{weight } t + w_a & \\ \quad \text{else} & \\ \quad \quad \text{weight } (\text{swapLeaves } t \ w_a \ a \ w_b \ b) = \text{weight } t & \\ \langle \text{proof} \rangle & \end{aligned}$$

lemma *cost_swapLeaves*:

$$\begin{aligned} \llbracket \text{consistent } t; a \neq b \rrbracket \implies & \\ \text{if } a \in \text{alphabet } t \text{ then} & \\ \quad \text{if } b \in \text{alphabet } t \text{ then} & \\ \quad \quad \text{cost } (\text{swapLeaves } t \ w_a \ a \ w_b \ b) + \text{freq } t \ a \times \text{depth } t \ a & \\ \quad \quad \quad + \text{freq } t \ b \times \text{depth } t \ b = & \\ \quad \quad \quad \text{cost } t + w_a \times \text{depth } t \ b + w_b \times \text{depth } t \ a & \\ \quad \text{else} & \\ \quad \quad \text{cost } (\text{swapLeaves } t \ w_a \ a \ w_b \ b) + \text{freq } t \ a \times \text{depth } t \ a = & \end{aligned}$$

$$\text{cost } t + w_b \times \text{depth } t a$$

else

if $b \in \text{alphabet } t$ then

$$\text{cost } (\text{swapLeaves } t w_a a w_b b) + \text{freq } t b \times \text{depth } t b =$$

$$\text{cost } t + w_a \times \text{depth } t b$$

else

$$\text{cost } (\text{swapLeaves } t w_a a w_b b) = \text{cost } t$$

$\langle \text{proof} \rangle$

Common sense tells us that the following statement is valid: “If Astrid exchanges her house with Bernard’s neighbor, Bernard becomes Astrid’s new neighbor.” A similar property holds for binary trees.

lemma *sibling_swapLeaves_sibling* [simp]:

$$\llbracket \text{consistent } t; \text{sibling } t b \neq b; a \neq b \rrbracket \implies$$

$$\text{sibling } (\text{swapLeaves } t w_a a w_s (\text{sibling } t b)) a = b$$

$\langle \text{proof} \rangle$

4.3 Symbol Interchange

The *swapSyms* function provides a simpler interface to *swapLeaves*, with *freq t a* and *freq t b* in place of w_a and w_b . Most lemmas about *swapSyms* are directly adapted from the homologous results about *swapLeaves*.

definition *swapSyms* :: $\alpha \text{ tree} \Rightarrow \alpha \Rightarrow \alpha \Rightarrow \alpha \text{ tree}$ **where**

$$\text{swapSyms } t a b \equiv \text{swapLeaves } t (\text{freq } t a) a (\text{freq } t b) b$$

lemma *swapSyms_id* [simp]:

$$\text{consistent } t \implies \text{swapSyms } t a a = t$$

$\langle \text{proof} \rangle$

lemma *alphabet_swapSyms* [simp]:

$$\llbracket a \in \text{alphabet } t; b \in \text{alphabet } t \rrbracket \implies \text{alphabet } (\text{swapSyms } t a b) = \text{alphabet } t$$

$\langle \text{proof} \rangle$

lemma *consistent_swapSyms* [simp]:

$$\text{consistent } t \implies \text{consistent } (\text{swapSyms } t a b)$$

$\langle \text{proof} \rangle$

lemma *depth_swapSyms_neither* [simp]:

$$\llbracket \text{consistent } t; c \neq a; c \neq b \rrbracket \implies$$

$$\text{depth } (\text{swapSyms } t a b) c = \text{depth } t c$$

$\langle \text{proof} \rangle$

lemma *freq_swapSyms* [simp]:

$$\llbracket \text{consistent } t; a \in \text{alphabet } t; b \in \text{alphabet } t \rrbracket \implies$$

$freq (swapSyms t a b) = freq t$
 ⟨proof⟩

lemma *cost_swapSyms*:

assumes *consistent t a ∈ alphabet t b ∈ alphabet t*

shows $cost (swapSyms t a b) + freq t a \times depth t a + freq t b \times depth t b =$
 $cost t + freq t a \times depth t b + freq t b \times depth t a$

⟨proof⟩

If a 's frequency is lower than or equal to b 's, and a is higher up in the tree than b or at the same level, then interchanging a and b does not increase the tree's cost.

lemma *le_le_imp_sum_mult_le_sum_mult*:

$\llbracket i \leq j; m \leq (n::nat) \rrbracket \implies i \times n + j \times m \leq i \times m + j \times n$

⟨proof⟩

lemma *cost_swapSyms_le*:

assumes *consistent t a ∈ alphabet t b ∈ alphabet t freq t a ≤ freq t b*
 $depth t a \leq depth t b$

shows $cost (swapSyms t a b) \leq cost t$

⟨proof⟩

As stated earlier, "If Astrid exchanges her house with Bernard's neighbor, Bernard becomes Astrid's new neighbor."

lemma *sibling_swapSyms_sibling [simp]*:

$\llbracket consistent t; sibling t b \neq b; a \neq b \rrbracket \implies$
 $sibling (swapSyms t a (sibling t b)) a = b$

⟨proof⟩

"If Astrid exchanges her house with Bernard, Astrid becomes Bernard's old neighbor's new neighbor."

lemma *sibling_swapSyms_other_sibling [simp]*:

$\llbracket consistent t; sibling t b \neq a; sibling t b \neq b; a \neq b \rrbracket \implies$
 $sibling (swapSyms t a b) (sibling t b) = a$

⟨proof⟩

4.4 Four-Way Symbol Interchange

The *swapSyms* function exchanges two symbols a and b . We use it to define the four-way symbol interchange function *swapFourSyms*, which takes four symbols a, b, c, d with $a \neq b$ and $c \neq d$, and exchanges them so that a and b occupy c and d 's positions.

A naive definition of this function would be

$$swapFourSyms t a b c d \equiv swapSyms (swapSyms t a c) b d.$$

This definition fails in the face of aliasing: If $a = d$, but $b \neq c$, then $\text{swapFourSyms } a b c d$ would leave a in b 's position. Incidentally, Cormen et al. [3, p. 390] forgot to consider this case in their proof.²

definition $\text{swapFourSyms} :: \alpha \text{ tree} \Rightarrow \alpha \Rightarrow \alpha \Rightarrow \alpha \Rightarrow \alpha \Rightarrow \alpha \text{ tree}$ **where**
 $\text{swapFourSyms } t a b c d \equiv$
 if $a = d$ then $\text{swapSyms } t b c$
 else if $b = c$ then $\text{swapSyms } t a d$
 else $\text{swapSyms } (\text{swapSyms } t a c) b d$

Lemmas about swapFourSyms are easy to prove by expanding its definition.

lemma $\text{alphabet_swapFourSyms}$ [simp]:
 $\llbracket a \in \text{alphabet } t; b \in \text{alphabet } t; c \in \text{alphabet } t; d \in \text{alphabet } t \rrbracket \implies$
 $\text{alphabet } (\text{swapFourSyms } t a b c d) = \text{alphabet } t$
 <proof>

lemma $\text{consistent_swapFourSyms}$ [simp]:
 $\text{consistent } t \implies \text{consistent } (\text{swapFourSyms } t a b c d)$
 <proof>

lemma freq_swapFourSyms [simp]:
 $\llbracket \text{consistent } t; a \in \text{alphabet } t; b \in \text{alphabet } t; c \in \text{alphabet } t;$
 $d \in \text{alphabet } t \rrbracket \implies$
 $\text{freq } (\text{swapFourSyms } t a b c d) = \text{freq } t$
 <proof>

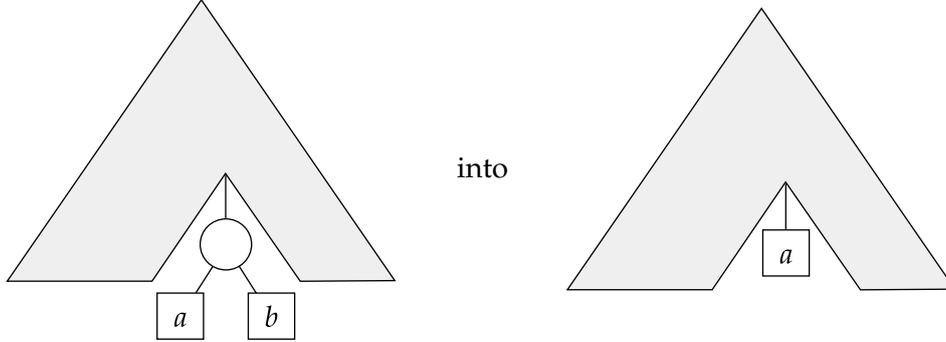
More Astrid and Bernard insanity: "If Astrid and Bernard exchange their houses with Carmen and her neighbor, Astrid and Bernard will now be neighbors."

lemma $\text{sibling_swapFourSyms_when_4th_is_sibling}$:
assumes $\text{consistent } t$ $a \in \text{alphabet } t$ $b \in \text{alphabet } t$ $c \in \text{alphabet } t$
 $a \neq b$ $\text{sibling } t c \neq c$
shows $\text{sibling } (\text{swapFourSyms } t a b c (\text{sibling } t c)) a = b$
 <proof>

²Thomas Cormen indicated in a personal communication that this will be corrected in the next edition of the book.

4.5 Sibling Merge

Given a symbol a , the *mergeSibling* function transforms the tree



The frequency of a in the result is the sum of the original frequencies of a and b , so as not to alter the tree's weight.

```
fun mergeSibling ::  $\alpha$  tree  $\Rightarrow$   $\alpha$   $\Rightarrow$   $\alpha$  tree where
mergeSibling (Leaf  $w_b$   $b$ )  $a$  = Leaf  $w_b$   $b$ 
mergeSibling (InnerNode  $w$  (Leaf  $w_b$   $b$ ) (Leaf  $w_c$   $c$ ))  $a$  =
  (if  $a = b \vee a = c$  then Leaf ( $w_b + w_c$ )  $a$ 
   else InnerNode  $w$  (Leaf  $w_b$   $b$ ) (Leaf  $w_c$   $c$ ))
mergeSibling (InnerNode  $w$   $t_1$   $t_2$ )  $a$  =
  InnerNode  $w$  (mergeSibling  $t_1$   $a$ ) (mergeSibling  $t_2$   $a$ )
```

The definition of *mergeSibling* has essentially the same structure as that of *sibling*. As a result, the custom induction rule that we derived for *sibling* works equally well for reasoning about *mergeSibling*.

lemmas *mergeSibling_induct_consistent* = *sibling_induct_consistent*

The properties of *mergeSibling* echo those of *sibling*. Like with *sibling*, simplification rules are crucial.

lemma *notin_alphabet_imp_mergeSibling_id* [*simp*]:
 $a \notin \text{alphabet } t \implies \text{mergeSibling } t \ a = t$
 <proof>

lemma *height_gt_0_imp_mergeSibling_left* [*simp*]:
 $\text{height } t_1 > 0 \implies$
 $\text{mergeSibling (InnerNode } w \ t_1 \ t_2) \ a =$
 $\text{InnerNode } w \ (\text{mergeSibling } t_1 \ a) \ (\text{mergeSibling } t_2 \ a)$
 <proof>

lemma *height_gt_0_imp_mergeSibling_right* [*simp*]:
 $\text{height } t_2 > 0 \implies$
 $\text{mergeSibling (InnerNode } w \ t_1 \ t_2) \ a =$

$InnerNode\ w\ (mergeSibling\ t_1\ a)\ (mergeSibling\ t_2\ a)$
 $\langle proof \rangle$

lemma *either_height_gt_0_imp_mergeSibling* [simp]:
 $height\ t_1 > 0 \vee height\ t_2 > 0 \implies$
 $mergeSibling\ (InnerNode\ w\ t_1\ t_2)\ a =$
 $InnerNode\ w\ (mergeSibling\ t_1\ a)\ (mergeSibling\ t_2\ a)$
 $\langle proof \rangle$

lemma *alphabet_mergeSibling* [simp]:
 $\llbracket consistent\ t; a \in alphabet\ t \rrbracket \implies$
 $alphabet\ (mergeSibling\ t\ a) = (alphabet\ t - \{sibling\ t\ a\}) \cup \{a\}$
 $\langle proof \rangle$

lemma *consistent_mergeSibling* [simp]:
 $consistent\ t \implies consistent\ (mergeSibling\ t\ a)$
 $\langle proof \rangle$

lemma *freq_mergeSibling*:
 $\llbracket consistent\ t; a \in alphabet\ t; sibling\ t\ a \neq a \rrbracket \implies$
 $freq\ (mergeSibling\ t\ a) =$
 $(\lambda c. \text{if } c = a \text{ then } freq\ t\ a + freq\ t\ (sibling\ t\ a)$
 $\quad \text{else if } c = sibling\ t\ a \text{ then } 0$
 $\quad \text{else } freq\ t\ c)$
 $\langle proof \rangle$

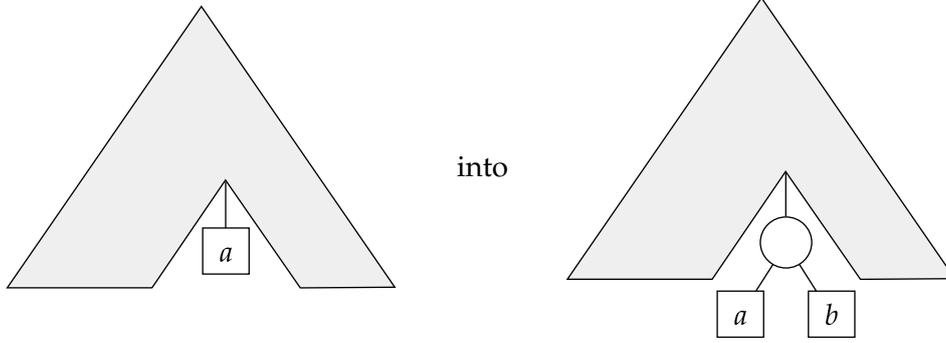
lemma *weight_mergeSibling* [simp]:
 $weight\ (mergeSibling\ t\ a) = weight\ t$
 $\langle proof \rangle$

If a has a sibling, merging a and its sibling reduces t 's cost by $freq\ t\ a + freq\ t\ (sibling\ t\ a)$.

lemma *cost_mergeSibling*:
 $\llbracket consistent\ t; sibling\ t\ a \neq a \rrbracket \implies$
 $cost\ (mergeSibling\ t\ a) + freq\ t\ a + freq\ t\ (sibling\ t\ a) = cost\ t$
 $\langle proof \rangle$

4.6 Leaf Split

The *splitLeaf* function undoes the merging performed by *mergeSibling*: Given two symbols a, b and two frequencies w_a, w_b , it transforms



In the resulting tree, a has frequency w_a and b has frequency w_b . We normally invoke it with w_a and w_b such that $\text{freq } t \ a = w_a + w_b$.

primrec *splitLeaf* :: $\alpha \text{ tree} \Rightarrow \text{nat} \Rightarrow \alpha \Rightarrow \text{nat} \Rightarrow \alpha \Rightarrow \alpha \text{ tree}$ **where**
splitLeaf (Leaf $w_c \ c$) $w_a \ a \ w_b \ b =$
 (if $c = a$ then InnerNode w_c (Leaf $w_a \ a$) (Leaf $w_b \ b$) else Leaf $w_c \ c$)
splitLeaf (InnerNode $w \ t_1 \ t_2$) $w_a \ a \ w_b \ b =$
 InnerNode w (*splitLeaf* $t_1 \ w_a \ a \ w_b \ b$) (*splitLeaf* $t_2 \ w_a \ a \ w_b \ b$)

primrec *splitLeaf_F* :: $\alpha \text{ forest} \Rightarrow \text{nat} \Rightarrow \alpha \Rightarrow \text{nat} \Rightarrow \alpha \Rightarrow \alpha \text{ forest}$ **where**
splitLeaf_F [] $w_a \ a \ w_b \ b = []$
splitLeaf_F ($t \cdot ts$) $w_a \ a \ w_b \ b =$
splitLeaf $t \ w_a \ a \ w_b \ b \cdot \text{splitLeaf_{F}} \ ts \ w_a \ a \ w_b \ b$

Splitting leaf nodes affects the alphabet, consistency, symbol frequencies, weight, and cost in unsurprising ways.

lemma *notin_alphabet_imp_splitLeaf_id* [simp]:
 $a \notin \text{alphabet } t \Longrightarrow \text{splitLeaf } t \ w_a \ a \ w_b \ b = t$
 <proof>

lemma *notin_alphabet_F_imp_splitLeaf_F_id* [simp]:
 $a \notin \text{alphabet}_F \ ts \Longrightarrow \text{splitLeaf}_F \ ts \ w_a \ a \ w_b \ b = ts$
 <proof>

lemma *alphabet_splitLeaf* [simp]:
 $\text{alphabet} (\text{splitLeaf } t \ w_a \ a \ w_b \ b) =$
 (if $a \in \text{alphabet } t$ then $\text{alphabet } t \cup \{b\}$ else $\text{alphabet } t$)
 <proof>

lemma *consistent_splitLeaf* [simp]:
 $[[\text{consistent } t; b \notin \text{alphabet } t]] \Longrightarrow \text{consistent} (\text{splitLeaf } t \ w_a \ a \ w_b \ b)$

$\langle \text{proof} \rangle$

lemma *freq_splitLeaf* [simp]:
[[consistent t; b \notin alphabet t]] \implies
freq (splitLeaf t w_a a w_b b) =
 (if a \in alphabet t then
 (λc . if c = a then w_a else if c = b then w_b else freq t c)
 else
 freq t)

$\langle \text{proof} \rangle$

lemma *weight_splitLeaf* [simp]:
[[consistent t; a \in alphabet t; freq t a = w_a + w_b]] \implies
weight (splitLeaf t w_a a w_b b) = weight t
 $\langle \text{proof} \rangle$

lemma *cost_splitLeaf* [simp]:
[[consistent t; a \in alphabet t; freq t a = w_a + w_b]] \implies
cost (splitLeaf t w_a a w_b b) = cost t + w_a + w_b
 $\langle \text{proof} \rangle$

4.7 Weight Sort Order

An invariant of Huffman's algorithm is that the forest is sorted by weight. This is expressed by the *sortedByWeight* function.

fun *sortedByWeight* :: α forest \Rightarrow bool **where**
sortedByWeight [] = True
sortedByWeight [t] = True
sortedByWeight (t₁ · t₂ · ts) =
 (weight t₁ \leq weight t₂ \wedge *sortedByWeight* (t₂ · ts))

The function obeys the following fairly obvious laws.

lemma *sortedByWeight_Cons_imp_sortedByWeight*:
sortedByWeight (t · ts) \implies *sortedByWeight* ts
 $\langle \text{proof} \rangle$

lemma *sortedByWeight_Cons_imp_forall_weight_ge*:
sortedByWeight (t · ts) \implies $\forall u \in \text{set } ts. \text{weight } u \geq \text{weight } t$
 $\langle \text{proof} \rangle$

lemma *sortedByWeight_insortTree*:
[[*sortedByWeight* ts; height t = 0; height_f ts = 0]] \implies
sortedByWeight (insortTree t ts)
 $\langle \text{proof} \rangle$

4.8 Pair of Minimal Symbols

The *minima* predicate expresses that two symbols $a, b \in \text{alphabet } t$ have the lowest frequencies in the tree t and that $\text{freq } t a \leq \text{freq } t b$. Minimal symbols need not be uniquely defined.

definition $\text{minima} :: \alpha \text{ tree} \Rightarrow \alpha \Rightarrow \alpha \Rightarrow \text{bool}$ **where**

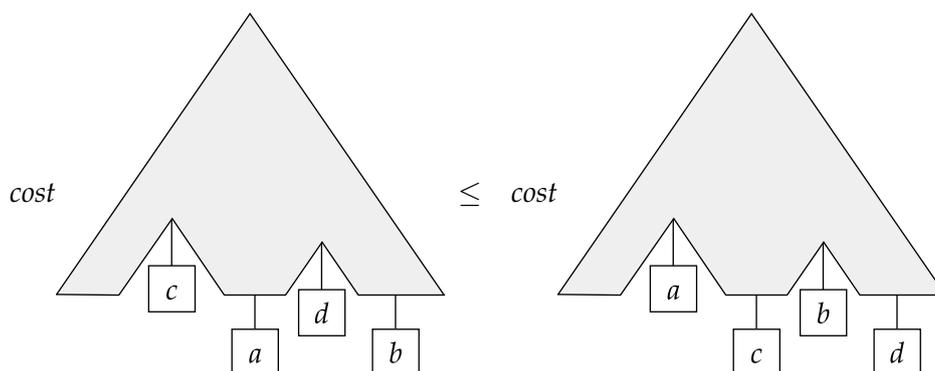
$\text{minima } t a b \equiv$

$$a \in \text{alphabet } t \wedge b \in \text{alphabet } t \wedge a \neq b \wedge \text{freq } t a \leq \text{freq } t b \\ \wedge (\forall c \in \text{alphabet } t. c \neq a \longrightarrow c \neq b \longrightarrow \\ \text{freq } t c \geq \text{freq } t a \wedge \text{freq } t c \geq \text{freq } t b)$$

5 Formalization of the Textbook Proof

5.1 Four-Way Symbol Interchange Cost Lemma

If a and b are minima, and c and d are at the very bottom of the tree, then exchanging a and b with c and d doesn't increase the cost. Graphically, we have



This cost property is part of Knuth's proof:

Let V be an internal node of maximum distance from the root. If w_1 and w_2 are not the weights already attached to the children of V , we can interchange them with the values that are already there; such an interchange does not increase the weighted path length.

Lemma 16.2 in Cormen et al. [3, p. 389] expresses a similar property, which turns out to be a corollary of our cost property:

Let C be an alphabet in which each character $c \in C$ has frequency $f[c]$. Let x and y be two characters in C having the lowest frequencies. Then there exists an optimal prefix code for C in which the codewords for x and y have the same length and differ only in the last bit.

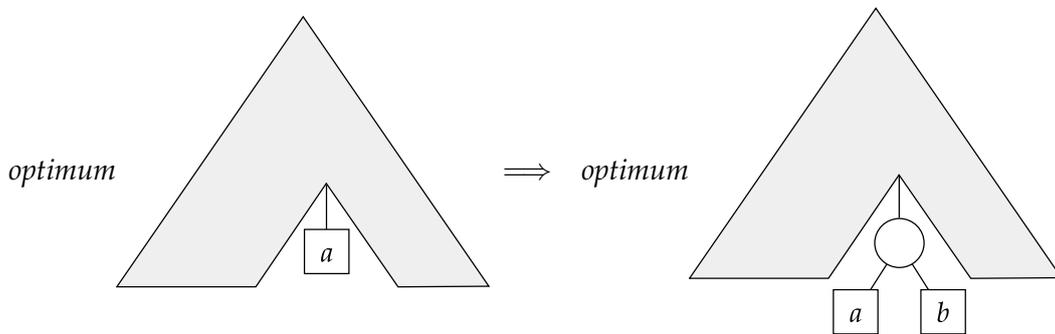
lemma *cost_swapFourSyms_le*:

assumes *consistent t minima t a b c ∈ alphabet t d ∈ alphabet t*
depth t c = height t depth t d = height t c ≠ d

shows *cost (swapFourSyms t a b c d) ≤ cost t*
 ⟨proof⟩

5.2 Leaf Split Optimality Lemma

The tree *splitLeaf t w_a a w_b b* is optimum if *t* is optimum, under a few assumptions, notably that *a* and *b* are minima of the new tree and that *freq t a = w_a + w_b*. Graphically:



This corresponds to the following fragment of Knuth’s proof:

Now it is easy to prove that the weighted path length of such a tree is minimized if and only if the tree with

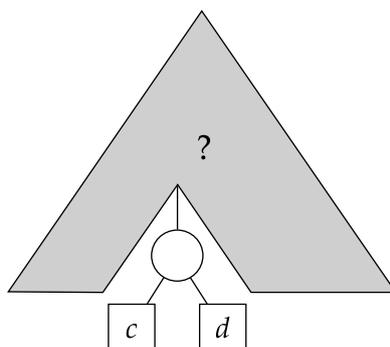


has minimum path length for the weights $w_1 + w_2, w_3, \dots, w_m$.

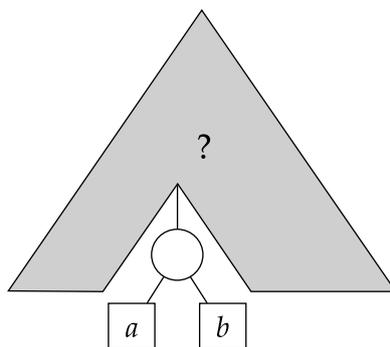
(We only need the “if” direction of Knuth’s equivalence.) Lemma 16.3 in Cormen et al. [3, p. 391] expresses essentially the same property:

Let C be a given alphabet with frequency $f[c]$ defined for each character $c \in C$. Let x and y be two characters in C with minimum frequency. Let C' be the alphabet C with characters x, y removed and (new) character z added, so that $C' = C - \{x, y\} \cup \{z\}$; define f for C' as for C , except that $f[z] = f[x] + f[y]$. Let T' be any tree representing an optimal prefix code for the alphabet C' . Then the tree T , obtained from T' by replacing the leaf node for z with an internal node having x and y as children, represents an optimal prefix code for the alphabet C .

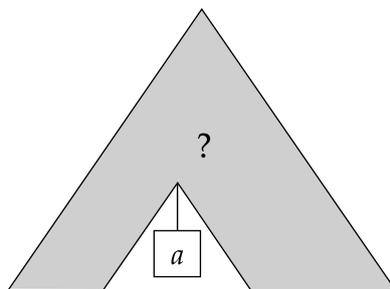
The proof is as follows: We assume that t has a cost less than or equal to that of any other compatible tree v and show that $\text{splitLeaf } t \ w_a \ a \ w_b \ b$ has a cost less than or equal to that of any other compatible tree u . By exists_at_height and $\text{depth_height_imp_sibling_ne}$, we know that some symbols c and d appear in sibling nodes at the very bottom of u :



(The question mark is there to remind us that we know nothing specific about u 's structure.) From u we construct a new tree $\text{swapFourSyms } u \ a \ b \ c \ d$ in which the minima a and b are siblings:



Merging a and b gives a tree compatible with t , which we can use to instantiate v in the assumption:



With this instantiation, the proof is easy:

$$\begin{aligned}
& \text{cost } (\text{splitLeaf } t \ a \ w_a \ b \ w_b) \\
= & \hspace{20em} (\text{cost_splitLeaf}) \\
& \text{cost } t + w_a + w_b \\
\leq & \hspace{20em} (\text{assumption}) \\
& \text{cost } (\overbrace{\text{mergeSibling } (\text{swapFourSyms } u \ a \ b \ c \ d) \ a}^v) + w_a + w_b \\
= & \hspace{15em} (\text{cost_mergeSibling}) \\
& \text{cost } (\text{swapFourSyms } u \ a \ b \ c \ d) \\
\leq & \hspace{15em} (\text{cost_swapFourSyms_le}) \\
& \text{cost } u.
\end{aligned}$$

The proof in Cormen et al. is by contradiction: Essentially, they assume that there exists a tree u with a lower cost than $\text{splitLeaf } t \ a \ w_a \ b \ w_b$ and show that there exists a tree v with a lower cost than t , contradicting the hypothesis that t is optimum. In place of $\text{cost_swapFourSyms_le}$, they invoke their lemma 16.2, which is questionable since u is not necessarily optimum.³

Our proof relies on the following lemma, which asserts that a and b are minima of u .

lemma *twice_freq_le_imp_minima*:

$\llbracket \forall c \in \text{alphabet } t. w_a \leq \text{freq } t \ c \wedge w_b \leq \text{freq } t \ c;$
 $\text{alphabet } u = \text{alphabet } t \cup \{b\}; a \in \text{alphabet } u; a \neq b;$
 $\text{freq } u = (\lambda c. \text{if } c = a \text{ then } w_a \text{ else if } c = b \text{ then } w_b \text{ else } \text{freq } t \ c);$
 $w_a \leq w_b \rrbracket \implies$
minima $u \ a \ b$
 $\langle \text{proof} \rangle$

Now comes the key lemma.

lemma *optimum_splitLeaf*:

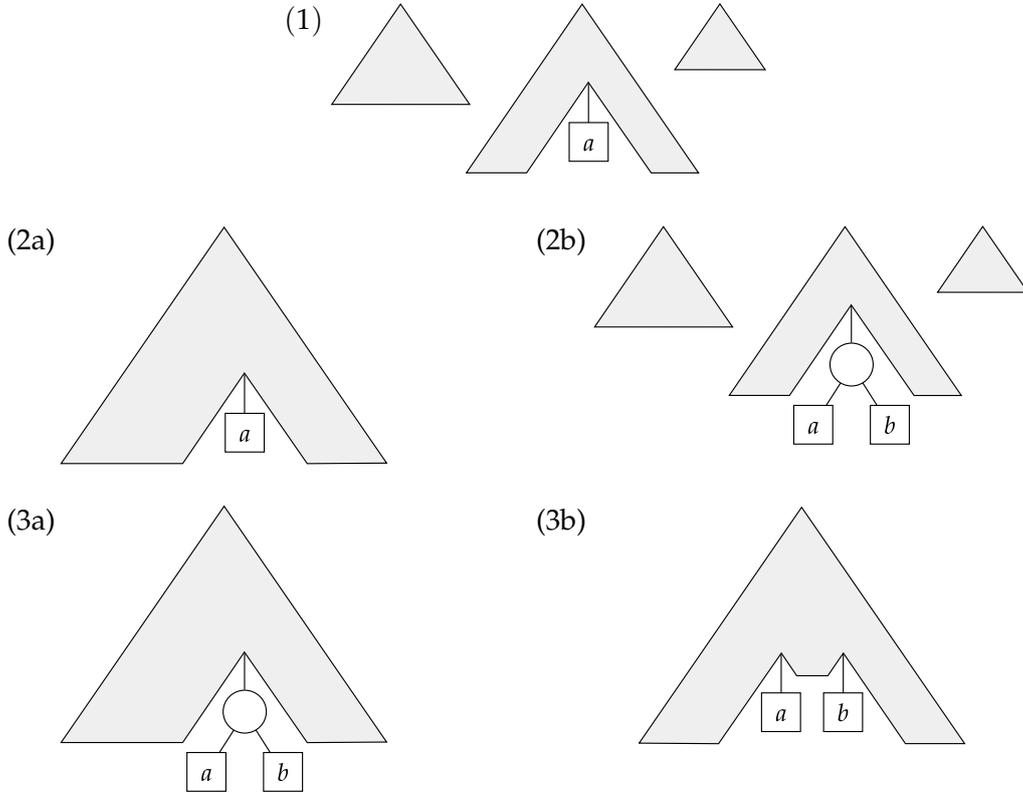
assumes *consistent* t *optimum* t $a \in \text{alphabet } t$ $b \notin \text{alphabet } t$
 $\text{freq } t \ a = w_a + w_b \quad \forall c \in \text{alphabet } t. \text{freq } t \ c \geq w_a \wedge \text{freq } t \ c \geq w_b$
 $w_a \leq w_b$

shows *optimum* $(\text{splitLeaf } t \ w_a \ a \ w_b \ b)$
 $\langle \text{proof} \rangle$

5.3 Leaf Split Commutativity Lemma

A key property of Huffman's algorithm is that once it has combined two lowest-weight trees using *uniteTrees*, it doesn't visit these trees ever again. This suggests that splitting a leaf node before applying the algorithm should give the same result as applying the algorithm first and splitting the leaf node afterward. The diagram below illustrates the situation:

³Thomas Cormen commented that this step will be clarified in the next edition of the book.



From the original forest (1), we can either run the algorithm (2a) and then split a (3a) or split a (2b) and then run the algorithm (3b). Our goal is to show that trees (3a) and (3b) are identical. Formally, we prove that

$$\text{splitLeaf} (\text{huffman } ts) w_a a w_b b = \text{huffman} (\text{splitLeaf}_F ts w_a a w_b b)$$

when ts is consistent, $a \in \text{alphabet}_F ts$, $b \notin \text{alphabet}_F ts$, and $\text{freq}_F ts a = w_a + w_b$. But before we can prove this commutativity lemma, we need to introduce a few simple lemmas.

lemma *cachedWeight_splitLeaf [simp]*:

$$\text{cachedWeight} (\text{splitLeaf } t w_a a w_b b) = \text{cachedWeight } t$$

<proof>

lemma *splitLeaf_F_insortTree_when_in_alphabet_left [simp]*:

$$\llbracket a \in \text{alphabet } t; \text{consistent } t; a \notin \text{alphabet}_F ts; \text{freq } t a = w_a + w_b \rrbracket \implies$$

$$\text{splitLeaf}_F (\text{insortTree } t ts) w_a a w_b b = \text{insortTree} (\text{splitLeaf } t w_a a w_b b) ts$$

<proof>

lemma *splitLeaf_F_insortTree_when_in_alphabet_tail [simp]*:

$$\llbracket a \in \text{alphabet}_F ts; \text{consistent}_F ts; a \notin \text{alphabet } t; \text{freq}_F ts a = w_a + w_b \rrbracket \implies$$

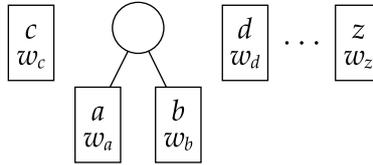
$$\text{splitLeaf}_F (\text{insortTree } t ts) w_a a w_b b =$$

$insertTree\ t\ (splitLeaf_{\mathbb{F}}\ ts\ w_a\ a\ w_b\ b)$
 $\langle proof \rangle$

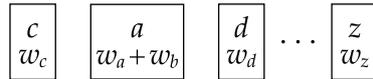
We are now ready to prove the commutativity lemma.

lemma $splitLeaf_huffman_commute$:
 $\llbracket consistent_{\mathbb{F}}\ ts; ts \neq []; a \in alphabet_{\mathbb{F}}\ ts; freq_{\mathbb{F}}\ ts\ a = w_a + w_b \rrbracket \implies$
 $splitLeaf\ (huffman\ ts)\ w_a\ a\ w_b\ b = huffman\ (splitLeaf_{\mathbb{F}}\ ts\ w_a\ a\ w_b\ b)$
 $\langle proof \rangle$

An important consequence of the commutativity lemma is that applying Huffman's algorithm on a forest of the form



gives the same result as applying the algorithm on the “flat” forest



followed by splitting the leaf node a into two nodes a, b with frequencies w_a, w_b . This effectively provides a way to flatten the forest at each step of the algorithm. Its invocation is implicit in the textbook proof.

5.4 Optimality Theorem

We are one lemma away from our main result.

lemma $max_0_imp_0$ [$simp$]:
 $(max\ x\ y = (0::nat)) = (x = 0 \wedge y = 0)$
 $\langle proof \rangle$

theorem $optimum_huffman$:
 $\llbracket consistent_{\mathbb{F}}\ ts; height_{\mathbb{F}}\ ts = 0; sortByWeight\ ts; ts \neq [] \rrbracket \implies$
 $optimum\ (huffman\ ts)$ $\langle proof \rangle$

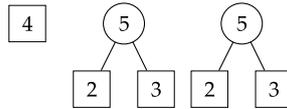
end

Theorem $optimum_huffman$ assumes that the forest ts passed to $huffman$ consists exclusively of leaf nodes. It is tempting to relax this restriction, by requiring instead that the forest ts is optimum. We would define optimality of a forest as

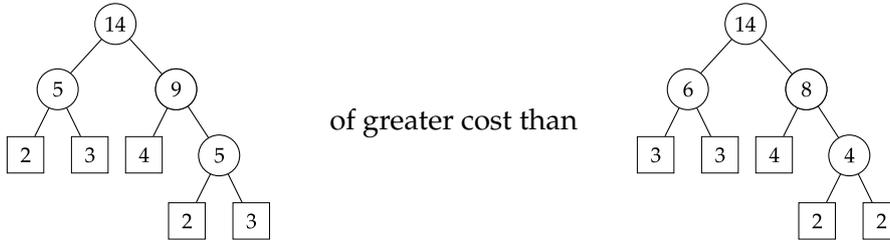
follows:

$$\begin{aligned} \text{optimum}_F ts \equiv & (\forall us. \text{length } ts = \text{length } us \longrightarrow \text{consistent}_F us \longrightarrow \\ & \text{alphabet}_F ts = \text{alphabet}_F us \longrightarrow \text{freq}_F ts = \text{freq}_F us \longrightarrow \\ & \text{cost}_F ts \leq \text{cost}_F us) \end{aligned}$$

with $\text{cost}_F [] = 0$ and $\text{cost}_F (t \cdot ts) = \text{cost } t + \text{cost}_F ts$. However, the modified proposition does not hold. A counterexample is the optimum forest



for which the algorithm constructs the tree



6 Related Work

Laurent Théry's Coq formalization of Huffman's algorithm [11, 12] is an obvious yardstick for our work. It has a somewhat wider scope, proving among others the isomorphism between prefix codes and binary trees. With 291 theorems, it is also much larger.

Théry identified the following difficulties in formalizing the textbook proof:

1. The leaf interchange process that brings the two minimal symbols together is tedious to formalize.
2. The sibling merging process requires introducing a new symbol for the merged node, which complicates the formalization.
3. The algorithm constructs the tree in a bottom-up fashion. While top-down procedures can usually be proved by structural induction, bottom-up procedures often require more sophisticated induction principles and larger invariants.
4. The informal proof relies on the notion of depth of a node. Defining this notion formally is problematic, because the depth can only be seen as a function if the tree is composed of distinct subtrees.

To circumvent these difficulties, Théry introduced the ingenious concept of cover. A forest ts is a *cover* of a tree t if t can be built from ts by adding inner nodes on top of the trees in ts . (The term “cover” is easier to understand if the binary trees are drawn with the root at the bottom of the page, like natural trees.) Huffman’s algorithm is then a refinement of the cover concept. The main proof consists in showing that the cost of *huffman* ts is less than or equal to that of any other tree for which ts is a cover. It relies on a few auxiliary definitions, notably an “ordered cover” concept that facilitates structural induction and a four-argument depth predicate (confusingly called *height*). Permutations also play a central role.

Incidentally, our experience suggests that the potential problems identified by Théry can be overcome without too much work:

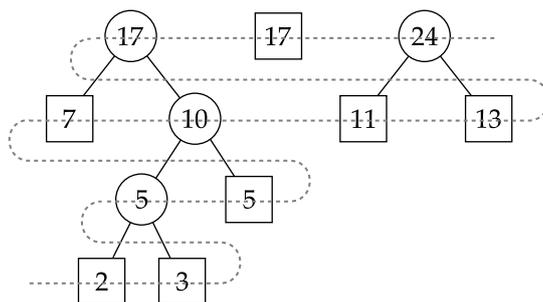
1. Formalizing the leaf interchange did not prove overly tedious. Among our 95 lemmas and theorems, 24 concern *swapLeaves*, *swapSyms*, and *swapFourSyms*.
2. The generation of a new symbol for the resulting node when merging two sibling nodes in *mergeSibling* was trivially solved by reusing one of the two merged symbols.
3. The bottom-up nature of the tree construction process was addressed by using the length of the forest as the induction measure and by merging the two minimal symbols, as in Knuth’s proof.
4. By restricting our attention to consistent trees, we were able to define the *depth* function simply and meaningfully.

7 Conclusion

The goal of most formal proofs is to increase our confidence in a result. In the case of Huffman’s algorithm, however, the chances that a bug would have gone unnoticed for the 56 years since its publication, under the scrutiny of leading computer scientists, seem extremely low; and the existence of a Coq proof should be sufficient to remove any remaining doubts.

The main contribution of this report was to demonstrate that the textbook proof of Huffman’s algorithm can be formalized in a straightforward manner using a state-of-the-art theorem prover such as Isabelle/HOL. In the process, we uncovered a few minor snags in the proof given in Cormen et al. [3]. Concerning Isabelle, the main lesson to draw from the Huffman proof is that custom induction rules, in combination with suitable simplification rules, greatly help the automatic proof tactics, sometimes reducing 30-line proof scripts to one-liners. We successfully applied this approach for handling both the ubiquitous “datatype + well-formedness predicate” combination (α *tree* + *consistent*) and functions defined by sequential pattern matching (*sibling* and *mergeSibling*).

In addition, formalizing the proof of Huffman’s algorithm led to a deeper understanding of this classic algorithm. Many of the lemmas, notably the leaf split commutativity lemma of Section 5.3, have not been found in the literature and express fundamental properties of the algorithm. Other discoveries didn’t find their way into the final proof. In particular, each step of the algorithm appears to preserve the invariant that the nodes in a forest are ordered by weight from left to right, bottom to top, as in the example below:



It is not hard to prove formally that a tree exhibiting this property is optimum. On the other hand, proving that the algorithm preserves this invariant seems difficult—more difficult than formalizing the textbook proof—and remains a suggestion for future work.

A few other directions for future work suggest themselves. First, we could formalize some of our hypotheses, notably our restriction to full and consistent binary trees. The next step could be to extend the proof’s scope to to cover *encode/decode* functions and connect prefix code trees to prefix codes, as done in the Coq development. Independently, we could generalize the development to n -ary trees.

Acknowledgments

I am grateful to several people for their help in producing this report. Tobias Nipkow suggested that I cut my teeth on Huffman coding and discussed several (sometimes flawed) drafts of the proof. He also provided many insights into Isabelle, which led to considerable simplifications. Alexander Krauss answered all my Isabelle questions and helped me with the trickier proofs. Thomas Cormen and Donald Knuth were both gracious enough to discuss their proofs with me, and Donald Knuth also suggested a terminology change. Finally, Mark Summerfield proposed many textual improvements.

References

- [1] Alfred V. Aho, John E. Hopcroft, and Jeffrey D. Ullman. *Data Structures and Algorithms*. Addison-Wesley, 1983.

- [2] Stephan Berghofer and Tobias Nipkow. Random testing in Isabelle/HOL. In J. Cuellar and Z. Liu, editors, *Software Engineering and Formal Methods (SEFM 2004)*, 230–239, IEEE Computer Society, 2004. Available online at <http://isabelle.in.tum.de/~nipkow/pubs/sefm04.html>.
- [3] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms* (Second Edition). MIT Press, 2001 and McGraw-Hill, 2002.
- [4] M. J. C. Gordon and T. F. Melham, editors. *Introduction to HOL: A Theorem Proving Environment for Higher Order Logic*. Cambridge University Press, 1993.
- [5] David A. Huffman. A method for the construction of minimum-redundancy codes. *Proceedings of the Institute of Radio Engineers* **40**(9):1098–1101, September 1952. Available online at http://compression.ru/download/articles/huff/huffman_1952_minimum-redundancy-codes.pdf.
- [6] Donald E. Knuth. *The Art of Computer Programming*, Vol. 1: *Fundamental Algorithms* (Third Edition). Addison-Wesley, 1997.
- [7] Alexander Krauss. *Defining Recursive Functions in Isabelle/HOL*. Department of Informatics, Technische Universität München, 2007. Updated version, <http://isabelle.in.tum.de/doc/functions.pdf>, June 8, 2008.
- [8] Robin Milner, Mads Tofte, Robert Harper, and David MacQueen. *The Definition of Standard ML* (Revised Edition). MIT Press, 1997.
- [9] Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL: A Proof Assistant for Higher-Order Logic*. Volume 2283 of *Lecture Notes in Computer Science*, Springer-Verlag, 2002. Updated version, <http://isabelle.in.tum.de/doc/tutorial.pdf>, June 8, 2008.
- [10] J. J. Rissanen. Generalized Kraft inequality and arithmetic coding. *IBM Journal of Research and Development* **20**(3):198–203, May 1976. Available online at <http://www.research.ibm.com/journal/rd/203/ibmrd2003B.pdf>.
- [11] Laurent Théry. *A Correctness Proof of Huffman Algorithm*. <http://coq.inria.fr/contribs/Huffman.html>, October 2003.
- [12] Laurent Théry. *Formalising Huffman’s Algorithm*. Technical report TRCS 034/2004, Department of Informatics, University of L’Aquila, 2004.
- [13] Markus Wenzel. *The Isabelle/Isar Reference Manual*. Department of Informatics, Technische Universität München, 2002. Updated version, <http://isabelle.in.tum.de/doc/isar-ref.pdf>, June 8, 2008.