Tobias Nipkow, Jasmin Blanchette,
Manuel Eberl, Alejandro Gómez-Londoño,
Peter Lammich, Christian Sternagel,
Simon Wimmer, Bohua Zhan

Functional Algorithms, Verified!

July 28, 2021

Preface

This book is an introduction to data structures and algorithms for functional languages, with a focus on proofs. It covers both functional correctness and running time analysis. It does so in a unified manner with inductive proofs about functional programs and their running time functions.

The unique feature of this book is that all the proofs have been machine-checked by the proof assistant Isabelle. That is, in addition to the text in the book, *which requires no knowledge of proof assistants*, there are the Isabelle definitions and proofs that can be accessed by following (in the PDF file) the links attached to section headings with a ↗ symbol. The structured nature of Isabelle proofs permits even novices to browse them and follow the high-level arguments.

This book has been classroom-tested for a number of years in a course for graduate and advanced undergraduate students. At the same time it is a reference for programmers and researchers who are interested in the details of some algorithm or proof.

Isabelle ↗

Isabelle [63, 76, 57] is a proof assistant for the logic HOL (= Higher-Order Logic), which is why the system is often called Isabelle/HOL. HOL is a generalization of first-order logic: functions can be passed as parameters and returned as results, just as in functional programming, and they can be quantified over. Isabelle's version of HOL also supports a simple version of Haskell's type classes.

The main strength of Isabelle and other proof assistants is their trustworthiness: all definitions, lemma statements, and inferences are checked. Beyond trustworthiness, formal proofs can also clarify arguments, by exposing and

explaining difficult steps. Most Isabelle users will confirm that their pen-and-paper proofs have become more lucid, and more correct, after they subjected themselves to the discipline of formal proof.

As emphasized above, the reader need not be familiar with Isabelle or HOL in order to read this book. However, to take full advantage of our proof assistant approach, readers are encouraged to learn how to write Isabelle definitions and proofs themselves — and to solve some of the exercises in this book. To this end we recommend the tutorial *Programming and Proving in Isabelle/HOL* [52], which is also Part I of the book *Concrete Semantics* [56].

Prerequisites

We expect the reader to be familiar with

- the basics of discrete mathematics: propositional and first-order logic, sets and relations, proof principles including induction;
- a typed functional programming language like Haskell [27], OCaml [60] or Standard ML [64];
- simple inductive proofs about functional programs.

Under Development

This book is meant to grow. New chapters are meant to be added over time. The list of authors is meant to grow — *you* could become one of them!

Colour

For the quick orientation of the reader, definitions are displayed in coloured boxes:

These boxes display functional programs.

These boxes display auxiliary definitions.

From a logical point of view there is no difference between the two kinds of definitions except that auxiliary definitions need not be executable.

Acknowledgements We are obviously indebted to the books by Cormen *et al.* [13] and Okasaki [62]. Mohammad Abdulaziz, Magnus Myreen and Larry Paulson commented on parts of the book. We are very grateful to all of them.

Contents

1 Basics 7

Part I Sorting and Selection

2 Sorting 19

3 Selection 35

Part II Search Trees

4 Binary Trees 49

5 Binary Search Trees 61

6 Abstract Data Types 75

7 2-3 Trees 81

8 Red-Black Trees 91

9 AVL Trees 101

10 Beyond Insert and Delete: $\cup$, $\cap$ and $-$ 113

11 Arrays via Braun Trees 121

12 Tries 143

13 Huffman's Algorithm 151

Part III Priority Queues

14 Priority Queues 167

15 Leftist Heaps 171

16 Priority Queues via Braun Trees 175

17 Binomial Heaps 179

Part IV Advanced Design and Analysis Techniques

18 Dynamic Programming 191

19 Amortized Analysis 213

20 Queues 217

21 Splay Trees 229

22 Skew Heaps 237

23 Pairing Heaps 241

Part V Appendix

A List Library 251

B Time Functions 255

C Notation 263

References 269

Index 275

1

Basics

In this chapter we describe the basic building blocks the book rests on.

Programs: The functional programming language we use is merely sketched because of its similarity with other well known functional languages.

Predefined types and notation: We introduce the basic predefined types and notations used in the book.

Inductive proofs: Although we do not explain proofs in general, we make an exception for certain inductive proofs.

Running time: We explain how we model running time by step counting functions.

1.1 Programs

The programs in this book are written in Isabelle's functional programming language which provides recursive algebraic data types (keyword: **datatype**), recursive function definitions and **let**, **if** and **case** expressions. The language is sufficiently close to a number of similar typed functional languages (SML [64], OCaml [60], Haskell [27]) to obviate the need for a detailed explanation. Moreover, Isabelle can generate SML, OCaml, Haskell and Scala code [24]. What distinguishes Isabelle's functional language from ordinary programming languages is that all functions in Isabelle must terminate. Termination must be proved. For all the functions in this book, termination is not difficult to see and Isabelle can prove it automatically. (If you want to go beyond, consult the function definition tutorial [44].)

Isabelle's functional language is pure logic. All language elements have precise definitions. However, this book is about algorithms, not programming language semantics. A functional programmer's intuition suffices for reading it. (If you want to know more about the logical basis see [7, 45, 26].)

A useful bit of notation: any infix operator can be turned into a function by enclosing it in parentheses, e.g. $(+)$.

1.2 Types

Type variables are denoted by $'a$, $'b$, etc. The function type arrow is $\Rightarrow$. Type constructor names follow their argument types, e.g. $'a\ list$. The notation $t :: \tau$ means that term t has type τ. The following types are predefined.

Booleans

Type *bool* comes with the constants *True* and *False* and the usual operations.

Numbers

There are three numeric types: the natural numbers *nat* $(0, 1, \ldots)$, the integers *int* and the real numbers *real*. They correspond to the mathematical sets $\mathbb{N}$, $\mathbb{Z}$ and $\mathbb{R}$ and not to any machine arithmetic. All three types come with the usual (overloaded) operations.

Sets

The type $'a\ set$ of sets (finite and infinite) over type $'a$ comes with the standard mathematical operations. The minus sign "$-$", unary or binary, can denote set complement or difference.

Lists

The type $'a\ list$ of lists whose elements are of type $'a$ is a recursive data type:

$$\textbf{datatype}\ 'a\ list = Nil \mid Cons\ 'a\ ('a\ list)$$

Constant *Nil* represents the empty list and $Cons\ x\ xs$ represents the list with first element x, the **head**, and rest list xs, the **tail**. The following syntactic sugar is sprinkled on top;

$$\begin{aligned} [] &\equiv Nil \\ x \# xs &\equiv Cons\ x\ xs \\ [x_1, \ldots, x_n] &\equiv x_1 \# \ldots \# x_n \# [] \end{aligned}$$

The $\equiv$ symbol means that the left-hand side is merely an abbreviation of the right-hand side.

A library of predefined functions on lists is shown in Appendix A. The length of a list xs is denoted by $|xs|$.

Type 'a option

The data type *'a option* is defined as follows:

datatype *'a option* = *None* | *Some 'a*

Pairs and Tuples

Pairs are written (a, b). Functions *fst* and *snd* select the first and second component of a pair: *fst* $(a, b) = a$ and *snd* $(a, b) = b$. The type *unit* contains only a single element $()$, the empty tuple.

1.2.1 Pattern Matching

Functions are defined by equations and pattern matching, for example over lists. Natural numbers may also be used in pattern-matching definitions:

fib $(n + 2) =$ *fib* $(n + 1) +$ *fib* n

Occasionally we use an extension of pattern matching where patterns can be named. For example, the defining equation

$f\ (x \# (y \# zs =: ys)) = ys$ @ zs

introduces a variables ys on the left that stands for $y \# zs$ and can be referred to on the right. Logically it is just an abbreviation of

$f\ (x \# y \# zs) =$ (**let** $ys = y \# zs$ **in** ys @ zs)

although it is suggestive of a more efficient interpretation. The general format is *pattern* =: *variable*.

1.2.2 Numeric Types and Coercions

The numeric types *nat*, *int* and *real* are all distinct. Converting between them requires explicit **coercion** functions, in particular the **inclusion** functions *int* :: *nat* ⇒ *int* and *real* :: *nat* ⇒ *real* that do not lose any information (in contrast to coercions in the other direction). We do not show inclusions unless they make a difference. For example, $(m + n)$ :: *real*, where m, n :: *nat*, is mathematically unambiguous because *real* $(m + n) =$ *real* $m +$ *real* n. On the other hand, $(m - n)$ :: *real* is ambiguous because *real* $(m - n) \neq$ *real* $m -$ *real* n because $(0::nat) - n = 0$. In some cases we can also drop coercions that are not inclusions, e.g. *nat* :: *int* ⇒ *nat*, which coerces negative integers to 0: if we know that $i \geq 0$ then we can drop the *nat* in *nat* i.

We prefer type *nat* over type *real* for ease of (Isabelle) proof. For example, for m, n :: *nat* we prefer $m \leq 2^n$ over lg $m \leq n$.

1.2.3 Multisets

Informally, a **multiset** is a set where elements can occur multiple times. Multisets come with the following operations:

$$
\begin{aligned}
\{\!\} &:: 'a\ multiset \\
(\in_\#) &:: 'a \Rightarrow 'a\ multiset \Rightarrow bool \\
add_mset &:: 'a \Rightarrow 'a\ multiset \Rightarrow 'a\ multiset \\
(+) &:: 'a\ multiset \Rightarrow 'a\ multiset \Rightarrow 'a\ multiset \\
size &:: 'a\ multiset \Rightarrow nat \\
mset &:: 'a\ list \Rightarrow 'a\ multiset \\
set_mset &:: 'a\ multiset \Rightarrow 'a\ set \\
image_mset &:: ('a \Rightarrow 'b) \Rightarrow 'a\ multiset \Rightarrow 'b\ multiset \\
filter_mset &:: ('a \Rightarrow bool) \Rightarrow 'a\ multiset \Rightarrow 'a\ multiset \\
sum_mset &:: 'a\ multiset \Rightarrow 'a
\end{aligned}
$$

Their meaning: $\{\!\}$ is the empty multiset; $(\in_\#)$ is the element test; *add_mset* adds an element to a multiset; $(+)$ is the sum of two multisets, where multiplicities of elements are added; *size* M, written $|M|$, is the number of elements in M, taking multiplicities into account; *mset* converts a list into a multiset by forgetting about the order of elements; *set_mset* converts a multiset into a set; *image_mset* applies a function to all elements of a multiset; *filter_mset* removes all elements from a multiset that do not satisfy the given predicate; *sum_mset* is the sum of the values of a multiset, the iteration of $(+)$ (taking multiplicity into account).

We use some additional suggestive syntax for some of these operations:

$$
\begin{aligned}
\{\!\{x \in_\# M \mid P\ x\}\!\} &\equiv filter_mset\ P\ M \\
\{\!\{f\ x \mid x \in_\# M\}\!\} &\equiv image_mset\ f\ M \\
\textstyle\sum_\# M &\equiv sum_mset\ M \\
\textstyle\sum_{x \in_\# M} f\ x &\equiv sum_mset\ (image_mset\ f\ M)
\end{aligned}
$$

See Section C.3 in the appendix for an overview of such syntax.

1.3 Notation

We adopt the following notation:

- The type of functions $'a \Rightarrow 'b$ comes with a predefined pointwise update operation with its own notation:
  $f(a := b) = (\lambda x.\ \textbf{if}\ x = a\ \textbf{then}\ b\ \textbf{else}\ f\ x)$
- Function lg is the binary logarithm.

We deviate from Isabelle's notation in favour of standard mathematics in a number of points:

- There is only one implication: $\Longrightarrow$ is printed as $\longrightarrow$ and $P \Longrightarrow Q \Longrightarrow R$ is printed as $P \wedge Q \longrightarrow R$.
- Multiplication is printed as $x \cdot y$.
- Exponentation is uniformly printed as x^y.
- We sweep under the carpet that type *nat* is defined as a recursive data type: **datatype** $nat = 0 \mid Suc\ nat$. In particular, constructor *Suc* is hidden: $Suc^k\ 0$ is printed as k and $Suc^k\ n$ (where n is not 0) is printed as $n + k$.
- Set comprehension syntax is the canonical $\{x \mid P\}$.

The reader who consults the Isabelle theories referred to in this book should be aware of these discrepancies.

1.4 Proofs

Proofs are the *raison d'être* of this book. Thus we present them in more detail than is customary in a book on algorithms. However, not all proofs:

- We omit proofs of simple properties of numbers, lists, sets and multisets, our predefined types. Obvious properties (e.g. $|xs\ @\ ys| = |xs| + |ys|$ or commutativity of $\cup$) are used implicitly without proof.
- With some exceptions, we only state properties if their proofs require induction, in which case we will say so, and we will always indicate which supporting properties were used.
- If a property is simply described as "inductive" or its proof is described by a phrase like "by an easy/automatic induction" it means that in the Isabelle proofs all cases of the induction were automatic, typically by simplification.

As a simple example of an easy induction consider the append function

$$
\begin{array}{l}
(@) :: \mathit{'a\ list} \Rightarrow \mathit{'a\ list} \Rightarrow \mathit{'a\ list} \\
[]\ @\ ys = ys \\
(x\ \#\ xs)\ @\ ys = x\ \#\ xs\ @\ ys
\end{array}
$$

and the proof of $(xs\ @\ ys)\ @\ zs = xs\ @\ ys\ @\ zs$ by structural induction on xs. (Note that $(@)$ associates to the right.) The base case is trivial by definition: $([]\ @\ ys)\ @\ zs = []\ @\ ys\ @\ zs$. The induction step is easy:

$$
\begin{array}{ll}
& (x\ \#\ xs\ @\ ys)\ @\ zs \\
= & x\ \#\ (xs\ @\ ys)\ @\ zs & \text{by definition of } (@) \\
= & x\ \#\ xs\ @\ ys\ @\ zs & \text{by IH}
\end{array}
$$

Note that **IH** stands for *Induction Hypothesis*, in this case $(xs \mathbin{@} ys) \mathbin{@} zs = xs \mathbin{@} ys \mathbin{@} zs$.

1.4.1 Computation Induction

Because most of our proofs are about recursive functions, most of them are by induction, and we say so explicitly. If we do not state explicitly what form the induction takes, it is by an obvious structural induction. The alternative and more general induction schema is **computation induction** where the induction follows the terminating computation, but from the bottom up. For example, the terminating recursive definition for $gcd :: nat \Rightarrow nat \Rightarrow nat$

$$gcd\ m\ n = (\textbf{if}\ n = 0\ \textbf{then}\ m\ \textbf{else}\ gcd\ n\ (m\ mod\ n))$$

gives rise to the following induction schema:

If $(n \neq 0 \longrightarrow P\ n\ (m\ mod\ n)) \longrightarrow P\ m\ n$ (for all m and n),
then $P\ m\ n$ (for all m and n).

In general, let $f :: \tau \Rightarrow \tau'$ be a terminating function of, for simplicity, one argument. Proving $P(x :: \tau)$ by induction on the computation of f means proving

$$P\ r_1 \wedge \ldots \wedge P\ r_n \longrightarrow P\ e$$

for every defining equation

$$f\ e = \ldots\ f\ r_1\ \ldots\ f\ r_n\ \ldots$$

where $f\ r_1, \ldots, f\ r_n$ are all the recursive calls. For simplicity we have ignored the **if** and **case** contexts that a recursive call $f\ r_i$ occurs in and that should be preconditions of the assumption $P\ r_i$ as in the *gcd* example. If the defining equations for f overlap, the above proof obligations are stronger than necessary.

1.5 Running Time

Our approach to reasoning about the **running time** of a function f is very simple: we explicitly define a function T_f such that $T_f\ x$ models the time the computation of $f\ x$ takes. Unless stated otherwise, T_f counts the number of all function calls in the computation of f. It is not intended that T_f yields the exact running time but only that the running time of f is in $O(T_f)$.

Given a function $f :: \tau_1 \Rightarrow \ldots \Rightarrow \tau_n \Rightarrow \tau$ we define a **(running) time function** $T_f :: \tau_1 \Rightarrow \ldots \Rightarrow \tau_n \Rightarrow nat$ by translating every defining equation for f into a defining equation for T_f. The translation is defined by two functions: $\mathcal{E}$ translates defining equations for f to defining equations for T_f and $\mathcal{T}$

translates expressions that compute some value to expressions that computes the number of function calls. The unusual notation $\mathcal{E}[\![.]\!]$ and $\mathcal{T}[\![.]\!]$ emphasizes that they are not functions in the logic.

$$\begin{array}{lcl} \mathcal{E}[\![f\ p_1\ \ldots\ p_n = e]\!] & = & (T_f\ p_1\ \ldots\ p_n = \mathcal{T}[\![e]\!] + 1) \\ \mathcal{T}[\![g\ e_1\ \ldots\ e_k]\!] & = & \mathcal{T}[\![e_1]\!] + \ldots + \mathcal{T}[\![e_k]\!] + T_g\ e_1\ \ldots\ e_k \end{array} \tag{1.1}$$

This is the general idea. It requires some remarks and clarifications:

- This definition of T_f is an abstraction of a call-by-value semantics. Thus it is also correct for lazy evaluation but may be a very loose upper bound.
- Definition (1.1) is incomplete: if g is a variable or constructor function (e.g. *Nil* or *Cons*), then there is no defining equation and thus no T_g. We simply define $T_g\ \ldots = 1$ if g is a variable, constructor function or predefined function on *bool* or numbers. These choices are discussed below.
- Conditionals and **case** expressions are treated specially:

$$\begin{array}{l} \mathcal{T}[\![\textbf{if } b \textbf{ then } e_1 \textbf{ else } e_2]\!] \\ = \mathcal{T}[\![b]\!] + (\textbf{if } b \textbf{ then } \mathcal{T}[\![e_1]\!] \textbf{ else } \mathcal{T}[\![e_2]\!]) \\ \mathcal{T}[\![\textbf{case } e \textbf{ of } p_1 \Rightarrow e_1 \mid \ldots \mid p_k \Rightarrow e_k]\!] \\ = \mathcal{T}[\![e]\!] + (\textbf{case } e \textbf{ of } p_1 \Rightarrow \mathcal{T}[\![e_1]\!] \mid \ldots \mid p_k \Rightarrow \mathcal{T}[\![e_k]\!]) \end{array}$$

- **let** expressions are also special:

$$\mathcal{T}[\![\textbf{let } x = e_1 \textbf{ in } e_2]\!] = \mathcal{T}[\![e_1]\!] + (\lambda x.\ \mathcal{T}[\![e_2]\!])\ e_1$$

  Note that $\mathcal{T}[\![e_2]\!]$ must be evaluated completely before $(\lambda x.\ \mathcal{T}[\![e_2]\!])\ e_1$ is β-contracted to ensure that $\mathcal{T}[\![x]\!] = 1$ (see above) kicks in instead of generating multiple copies of $\mathcal{T}[\![e_1]\!]$.
- For simplicity we restrict ourselves to a first-order language above. Nevertheless we use a few basic higher-order functions like *map* in the book. Their running time functions are defined in Appendix B.1.

As an example consider the append function (@) defined above. The defining equations for $T_{append} :: 'a\ list \Rightarrow 'a\ list \Rightarrow nat$ are easily derived. The first equation translates like this:

$$\begin{array}{l} \mathcal{E}[\![[] \ @\ ys = ys]\!] \\ = (T_{append}\ []\ ys = \mathcal{T}[\![ys]\!] + 1) \\ = (T_{append}\ []\ ys = 2) \end{array}$$

The right-hand side of the second equation translates like this:

$$\begin{aligned}&\mathcal{T}[\![x \mathbin{\#} xs \mathbin{@} ys]\!]\\ &= \mathcal{T}[\![x]\!] + \mathcal{T}[\![xs \mathbin{@} ys]\!] + T_{Cons}\ x\ (xs \mathbin{@} ys)\\ &= 1 + (\mathcal{T}[\![xs]\!] + \mathcal{T}[\![ys]\!] + T_{append}\ xs\ ys) + 1\\ &= 1 + (1 + 1 + T_{append}\ xs\ ys) + 1\end{aligned}$$

Thus the two defining equations for T_{append} are

$$\begin{aligned}&T_{append}\ []\ ys = 2\\ &T_{append}\ (x \mathbin{\#} xs)\ ys = T_{append}\ xs\ ys + 5\end{aligned}$$

The constants 2 and 5 in the above definition are somewhat arbitrary because accessing a variable and calling $(\#)$ take different amounts of time, even if we assume they are both constant-time operations. Because we are only interested in $O(T_f)$ it is justified to reduce all additive constants to 1. This is what we will call the **canonical running time function**. Thus the canonical T_{append} is this:

$$\begin{aligned}&T_{append}\ []\ ys = 1\\ &T_{append}\ (x \mathbin{\#} xs)\ ys = T_{append}\ xs\ ys + 1\end{aligned}$$

Unless noted otherwise we always work with canonical running time functions. This may also involve looking inside **if**/**case** when reducing additive constants, e.g. reducing (**if** b **then** $2 + e$ **else** $3) + 2$ to **if** b **then** $1 + e$ **else** 1.

Occasionally we will apply one more simplification step: when defining some T_f, you may drop additive constants when they are added to some term $T_g \ldots$ where T_g has already been defined (i.e. $g \neq f$ and f and g are not mutually recursive). This is because by construction all T_g return positive results and adding a constant does not change $O(T_f)$.

In the main body of the book we initially show the definition of each T_f. Once the principles above have been exemplified sufficiently, the time functions are relegated to Appendix B.

1.5.1 Example: List Reversal

This section exemplifies not just the definition of time functions but also their analysis. The standard list reversal function *rev* is defined in Appendix A. This is the corresponding canonical time function:

$$\begin{aligned}&T_{rev} :: 'a\ list \Rightarrow nat\\ &T_{rev}\ []\ = 1\\ &T_{rev}\ (x \mathbin{\#} xs) = T_{rev}\ xs + T_{append}\ (rev\ xs)\ [x] + 1\end{aligned}$$

The canonical T_{append} was shown before.

A simple induction shows $T_{append}\ xs\ ys = |xs| + 1$. The precise formula for T_{rev} is less immediately obvious (exercise!) but an upper bound is easy to guess and verify by induction:

$$T_{rev}\ xs \le (|xs| + 1)^2$$

We will frequently prove upper bounds only.

Of course one can also reverse a list in linear time:

```
itrev :: 'a list ⇒ 'a list ⇒ 'a list
itrev [] ys = ys
itrev (x # xs) ys = itrev xs (x # ys)
```

```
T_itrev :: 'a list ⇒ 'a list ⇒ nat
T_itrev [] _ = 1
T_itrev (x # xs) ys = T_itrev xs (x # ys) + 1
```

Function *itrev* has linear running time: $T_{itrev}\ xs\ ys = |xs| + 1$. A simple induction yields $itrev\ xs\ ys = rev\ xs\ @\ ys$. Thus *itrev* implements *rev*: $rev\ xs = itrev\ xs\ []$.

1.5.2 Discussion

Analysing the running time of any program, let alone a functional one, is tricky unless you have two formally specified entities: a machine model that includes a notion of running time and a translation from programs to machines. For imperative programs the standard model is the Random Access Machine (RAM) where each instruction takes one time unit. For functional programs a standard measure is the number of function calls. We adopt this measure and equate it with time although a more honest name would be computation steps.

A full proof that the execution time of our functional programs is in $O(T_f)$ on some actual soft- and hardware is a major undertaking: one would need to formalize the full stack of compiler, runtime system and hardware. We do not offer such a proof. Thus our formalization of "time" should be seen as conditional: given a stack that satisfies our basic assumptions in the definition of $\mathcal{E}$ and $\mathcal{T}$, our analyses are correct for that stack. We argue that these assumptions are not unreasonable (on a RAM) provided we accept that both the address space and numbers have a fixed size and cannot grow arbitrarily.

Of course this means that actual program execution may abort if the resources are exhausted.

Our basic assumption is that function calls take constant time. This is reasonable because we just need to allocate (and later deallocate) a new stack frame of constant size because all parameters are references or numbers and thus of fixed size. We also assumed that variable access takes constant time. This is a standard RAM assumption. Assuming that constructor functions take constant time is reasonable because the memory manager could simply employ a single reference to the first free memory cell and increment that with each constructor function call. How to account for garbage collection is less clear. In the worst case we have to assume that garbage collection is switched off, which simply exhausts memory more quickly. Finally we assume that operations on *bool* and numbers take constant time. The former is obvious, the latter follows from our assumption that we have fixed-size numbers.

In the end, we are less interested in a specific model of time and more in the principle that time (and other resources) can be analyzed just as formally as functional correctness once the ground rules (e.g. T) have been established.

1.5.3 Asymptotic Notation

The above approach to running time analysis is nicely concrete and avoids the more sophisticated machinery of asymptotic notation, $O(.)$ and friends. Thus we have intentionally lowered the entry barrier to the book for readers who want to follow the Isabelle formalization: we require no familiarity with Isabelle's real analysis library and in particular with the existing formalization of and automation for asymptotic notation [17]. Of course this comes at a price: one has to come up with and reason about somewhat arbitrary constants in the analysis of individual functions. Moreover we rarely appeal to the "master theorem" (although Eberl [17] provides a generalized version) but prove solutions to recurrence equations correct by induction. Again, this is merely to reduce the required mathematical basis and to show that it can be done. In informal explanations, typically when considering inessential variations, we do use standard mathematical notation and write, for example, $O(n \lg n)$.

Part I

Sorting and Selection

2

Sorting

In this chapter we define and verify the following sorting functions: insertion sort, quicksort, and one top-down and two bottom-up merge sorts. We also analyze their running times, except for quicksort.

Sorting involves an ordering. In this book we assume such an ordering to be provided by comparison operators $\leq$ and $<$ defined on the underlying type.

Sortedness of lists is defined as follows:

```
sorted :: ('a::linorder) list ⇒ bool
sorted [] = True
sorted (x # ys) = ((∀ y∈set ys. x ≤ y) ∧ sorted ys)
```

That is, every element is $\leq$ than all elements to the right of it: the list is sorted in increasing order. But what does $'a::linorder$ mean?

The type variable $'a$ is annotated with *linorder*, which means that *sorted* is only applicable if a binary predicate $(\leq) :: 'a \Rightarrow 'a \Rightarrow bool$ is defined and $(\leq)$ is a **linear order**, i.e. the following properties are satisfied:

reflexivity:	$x \leq x$
transitivity:	$x \leq y \wedge y \leq z \longrightarrow x \leq z$
antisymmetry:	$a \leq b \wedge b \leq a \longrightarrow a = b$
linearity/totality:	$x \leq y \vee y \leq x$

Moreover, the binary predicate $(<)$ must satisfy

$$x < y \longleftrightarrow x \leq y \wedge x \neq y.$$

Note that *linorder* is a specific predefined example of a **type class** [25]. We will not explain type classes any further because we do not require the general concept. In fact, we will mostly not even show the *linorder* annotation

in types: you can assume that if you see $\leq$ or $<$ on a generic type $'a$ in this book, $'a$ is implicitly annotated with *linorder*. Note further that $(\leq)$ on the numeric types *nat*, *int* and *real* is a linear order.

2.1 Specification of Sorting Functions

A sorting function $sort :: 'a\ list \Rightarrow 'a\ list$ (where, as usual, $'a :: linorder$) must obviously satisfy the following property:

$$sorted\ (sort\ xs)$$

However, this is not enough — otherwise, $wrong_sort\ xs = []$ would be a correct sorting function. The set of elements in the output must be the same as in the input, and each element must occur the same number of times. This is most readily captured with the notion of a multiset (see Section 1.2.3). The second property that a sorting function *sort* must satisfy is

$$mset\ (sort\ xs) = mset\ xs$$

where function *mset* converts a list into its corresponding multiset.

2.2 Insertion Sort

Insertion sort is well-known for its intellectual simplicity and computational inefficiency. Its simplicity makes it an ideal starting point for this book. Below, it is implemented by the function *isort* with the help of the auxiliary function *insort* that inserts a single element into an already sorted list.

```
insort :: 'a ⇒ 'a list ⇒ 'a list
insort x [] = [x]
insort x (y # ys) = (if x ≤ y then x # y # ys else y # insort x ys)

isort :: 'a list ⇒ 'a list
isort [] = []
isort (x # xs) = insort x (isort xs)
```

2.2.1 Functional Correctness

We start by proving the preservation of the multiset of elements:

$$mset\ (insort\ x\ xs) = \{\!\{x\}\!\} + mset\ xs \tag{2.1}$$

$$mset\ (isort\ xs) = mset\ xs \tag{2.2}$$

Both properties are proved by induction; the proof of (2.2) requires (2.1).

Now we turn to sortedness. Because the definition of *sorted* involves *set*, it is frequently helpful to prove multiset preservation first (as we have done above) because that yields preservation of the set of elements. That is, from (2.1) we obtain:

$$set\ (insort\ x\ xs) = \{x\} \cup set\ xs \tag{2.3}$$

Two inductions prove

$$sorted\ (insort\ a\ xs) = sorted\ xs \tag{2.4}$$

$$sorted\ (isort\ xs) \tag{2.5}$$

where the proof of (2.4) uses (2.3) and the proof of (2.5) uses (2.4).

2.2.2 Running Time Analysis

These are the canonical running time functions (according to Section 1.5):

$T_{insort} :: 'a \Rightarrow 'a\ list \Rightarrow nat$
$T_{insort}\ _\ [] = 1$
$T_{insort}\ x\ (y\ \#\ ys) = (\textbf{if}\ x \le y\ \textbf{then}\ 0\ \textbf{else}\ T_{insort}\ x\ ys) + 1$

$T_{isort} :: 'a\ list \Rightarrow nat$
$T_{isort}\ [] = 1$
$T_{isort}\ (x\ \#\ xs) = T_{isort}\ xs + T_{insort}\ x\ (isort\ xs) + 1$

A dismal quadratic upper bound for the running time of insertion sort is proved readily:

Lemma 2.1. $T_{isort}\ xs \le (|xs| + 1)^2$

Proof. The following properties are proved by induction on xs:

$$T_{insort}\ x\ xs \le |xs| + 1 \tag{2.6}$$

$$|insort\ x\ xs| = |xs| + 1 \tag{2.7}$$

$$|isort\ xs| = |xs| \tag{2.8}$$

The proof of (2.8) needs (2.7). The proof of $T_{isort}\ xs \le (|xs| + 1)^2$ is also by induction on xs. The base case is trivial. The induction step is easy:

$$\begin{aligned}
&T_{isort}\ (x\ \#\ xs) = T_{isort}\ xs + T_{insort}\ x\ (isort\ xs) + 1 \\
&\leq (|xs| + 1)^2 + T_{insort}\ x\ (isort\ xs) + 1 && \text{by IH} \\
&\leq (|xs| + 1)^2 + |xs| + 1 + 1 && \text{using (2.6) and (2.8)} \\
&\leq (|x\ \#\ xs| + 1)^2 && \square
\end{aligned}$$

Exercise 2.2 asks you to show that *isort* actually has quadratic running time on all lists $[n, n-1, \ldots, 0]$.

2.2.3 Exercises

Exercise 2.1. Show that any sorting function has the same input/output behaviour as insertion sort:

$$(\forall xs.\ mset\ (f\ xs) = mset\ xs) \wedge (\forall xs.\ sorted\ (f\ xs)) \longrightarrow f\ xs = isort\ xs$$

Exercise 2.2. Show that T_{isort} achieves its optimal value of $2 \cdot n + 1$ for sorted lists, and its worst-case value of $(n + 1) \cdot (n + 2)\ div\ 2$ for the list $rev\ [0..<n]$.

2.3 Quicksort

Quicksort is a divide-and-conquer algorithm that sorts a list as follows: pick a **pivot** element from the list; partition the remaining list into those elements that are smaller and those that are greater than the pivot (equal elements can go into either sublist); sort these sublists recursively and append the results. A particularly simple version of this approach, where the first element is chosen as the pivot, looks like this:

```
quicksort :: 'a list ⇒ 'a list

quicksort [] = []
quicksort (x # xs)
= quicksort (filter (λy. y < x) xs) @
  [x] @ quicksort (filter (λy. y ≥ x) xs)
```

2.3.1 Functional Correctness

Preservation of the multiset of elements

$$mset\ (quicksort\ xs) = mset\ xs \tag{2.9}$$

is proved by computation induction using these lemmas:

$$mset\ (filter\ P\ xs) = filter_mset\ P\ (mset\ xs)$$
$$(\forall x.\ P\ x = (\neg\ Q\ x)) \longrightarrow filter_mset\ P\ M\ +\ filter_mset\ Q\ M\ =\ M$$

A second computation induction proves sortedness

$$sorted\ (quicksort\ xs)$$

using the lemmas

$$sorted\ (xs\ @\ ys)$$
$$= (sorted\ xs \wedge sorted\ ys \wedge (\forall x{\in}set\ xs.\ \forall y{\in}set\ ys.\ x \leq y))$$
$$set\ (quicksort\ xs) = set\ xs$$

where the latter one is an easy consequence of (2.9).

We do not analyze the running time of *quicksort*. It is well known that in the worst case it is quadratic (exercise!) but that the average-case running time (in a certain sense) is $O(n \lg n)$. If the pivot is chosen randomly instead of always choosing the first element, the *expected* running time is also $O(n \lg n)$. The necessary probabilistic analysis is beyond the scope of this text but can be found elsewhere [16, 18].

2.3.2 Exercises

Exercise 2.3. Function *quicksort* appends the lists returned from the recursive calls. This is expensive because the running time of (@) is linear in the length of its first argument. Define a function *quicksort2* :: *'a list* ⇒ *'a list* ⇒ *'a list* that avoids (@) but accumulates the result in its second parameter via (#) only. Prove *quicksort2 xs ys* = *quicksort xs* @ *ys*.

Exercise 2.4. There is one obvious optimisation to the version of quicksort that we studied before: instead of partitioning the list into those elements that are smaller than the pivot and those that are at least as big as the pivot, we can use three-way partitioning:

```
partition3 :: 'a ⇒ 'a list ⇒ 'a list × 'a list × 'a list
partition3 x xs
= (filter (λy. y < x) xs, filter (λy. y = x) xs,
   filter (λy. y > x) xs)

quicksort3 :: 'a list ⇒ 'a list
quicksort3 [] = []
```

$$
\begin{aligned}
&\mathit{quicksort3}\ (x\ \#\ xs)\\
&= (\textbf{let}\ (ls,\ es,\ gs) = \mathit{partition3}\ x\ xs\\
&\quad\ \textbf{in}\ \mathit{quicksort3}\ ls\ @\ x\ \#\ es\ @\ \mathit{quicksort3}\ gs)
\end{aligned}
$$

Prove that this version of quicksort also produces the correct results.

Exercise 2.5. In this exercise, we will examine the worst-case behaviour of Quicksort, which is e.g. achieved if the input list is already sorted. Consider the time function for Quicksort:

$$
\begin{aligned}
&T_{\mathit{quicksort}} :: \mathit{'a\ list} \Rightarrow \mathit{nat}\\
&T_{\mathit{quicksort}}\ [] = 1\\
&T_{\mathit{quicksort}}\ (x\ \#\ xs) = T_{\mathit{quicksort}}\ (\mathit{filter}\ (\lambda y.\ y < x)\ xs) +\\
&\qquad\qquad\qquad\qquad T_{\mathit{quicksort}}\ (\mathit{filter}\ (\lambda y.\ y \geq x)\ xs) +\\
&\qquad\qquad\qquad\qquad 2 \cdot T_{\mathit{filter}}\ (\lambda_.\ 1)\ xs + 1
\end{aligned}
$$

1. Show that Quicksort takes quadratic time on sorted lists by proving

$$\mathit{sorted}\ xs \longrightarrow T_{\mathit{quicksort}}\ xs = a \cdot |xs|^2 + b \cdot |xs| + c$$

   for suitable values a, b, c.
2. Show that this is the worst-case running time by proving

$$T_{\mathit{quicksort}}\ xs \leq a \cdot |xs|^2 + b \cdot |xs| + c$$

   for the values of a, b, c you determined in the previous step.

2.4 Top-Down Merge Sort

Merge sort is another prime example of a divide-and-conquer algorithm, and one whose running time is guaranteed to be $O(n \lg n)$. We will consider three variants and start with the simplest one: split the list into two halves, sort the halves separately and merge the results.

$$
\begin{aligned}
&\mathit{merge} :: \mathit{'a\ list} \Rightarrow \mathit{'a\ list} \Rightarrow \mathit{'a\ list}\\
&\mathit{merge}\ []\ ys = ys
\end{aligned}
$$

```
merge xs [] = xs
merge (x # xs) (y # ys)
= (if x ≤ y then x # merge xs (y # ys)
   else y # merge (x # xs) ys)

msort :: 'a list ⇒ 'a list
msort xs
= (let n = |xs|
   in if n ≤ 1 then xs
      else merge (msort (take (n div 2) xs))
                 (msort (drop (n div 2) xs)))
```

2.4.1 Functional Correctness

We start off with multisets and sets of elements:

$$mset\ (merge\ xs\ ys) = mset\ xs + mset\ ys \tag{2.10}$$

$$set\ (merge\ xs\ ys) = set\ xs \cup set\ ys \tag{2.11}$$

Proposition (2.10) is proved by induction on the computation of *merge* and (2.11) is an easy consequence.

Lemma 2.2. $mset\ (msort\ xs) = mset\ xs$

Proof. The proof is by induction on the computation of *msort*. Let $n = |xs|$. The base case ($n \leq 1$) is trivial. Now assume $n > 1$ and let $ys = take\ (n\ div\ 2)\ xs$ and $zs = drop\ (n\ div\ 2)\ xs$.

$mset\ (msort\ xs) = mset\ (msort\ ys) + mset\ (msort\ zs)$ by (2.10)
$= mset\ ys + mset\ zs$ by IH
$= mset\ (ys\ @\ zs)$
$= mset\ xs$ □

Now we turn to sortedness. An induction on the computation of *merge*, using (2.11), yields

$$sorted\ (merge\ xs\ ys) = (sorted\ xs \wedge sorted\ ys) \tag{2.12}$$

Lemma 2.3. $sorted\ (msort\ xs)$

The proof is an easy induction on the computation of *msort*. The base case ($n \leq 1$) follows because every list of length ≤ 1 is sorted. The induction step follows with the help of (2.12).

2.4.2 Running Time Analysis

To simplify the analysis, and in line with the literature, we only count the number of comparisons:

$$
\begin{array}{l}
C_{merge} :: 'a\ list \Rightarrow 'a\ list \Rightarrow nat \\
C_{merge}\ []\ _ = 0 \\
C_{merge}\ _\ []\ = 0 \\
C_{merge}\ (x\ \#\ xs)\ (y\ \#\ ys) \\
= 1 + (\textbf{if}\ x \le y\ \textbf{then}\ C_{merge}\ xs\ (y\ \#\ ys)\ \textbf{else}\ C_{merge}\ (x\ \#\ xs)\ ys)
\end{array}
$$

$$
\begin{array}{l}
C_{msort} :: 'a\ list \Rightarrow nat \\
C_{msort}\ xs \\
= (\textbf{let}\ n = |xs|; \\
\qquad ys = take\ (n\ div\ 2)\ xs; \\
\qquad zs = drop\ (n\ div\ 2)\ xs \\
\quad \textbf{in if}\ n \le 1\ \textbf{then}\ 0 \\
\qquad \textbf{else}\ C_{msort}\ ys + C_{msort}\ zs + C_{merge}\ (msort\ ys)\ (msort\ zs))
\end{array}
$$

By computation inductions we obtain:

$$|merge\ xs\ ys| = |xs| + |ys| \tag{2.13}$$

$$|msort\ xs| = |xs| \tag{2.14}$$

$$C_{merge}\ xs\ ys \le |xs| + |ys| \tag{2.15}$$

where the proof of (2.14) uses (2.13).

To simplify technicalities, we prove the $n \cdot \lg n$ bound on the number of comparisons in *msort* only for $n = 2^k$, in which case the bound becomes $k \cdot 2^k$.

Lemma 2.4. $|xs| = 2^k \longrightarrow C_{msort}\ xs \le k \cdot 2^k$

Proof. The proof is by induction on k. The base case is trivial and we concentrate on the step. Let $n = |xs|$, $ys = take\ (n\ div\ 2)\ xs$ and $zs = drop\ (n\ div\ 2)\ xs$. The case $n \le 1$ is trivial. Now assume $n > 1$.

$$
\begin{array}{lll}
& C_{msort}\ xs & \\
= & C_{msort}\ ys + C_{msort}\ zs + C_{merge}\ (msort\ ys)\ (msort\ zs) & \\
\le & C_{msort}\ ys + C_{msort}\ zs + |ys| + |zs| & \text{using (2.15) and (2.14)} \\
\le & k \cdot 2^k + k \cdot 2^k + |ys| + |zs| & \text{by IH} \\
= & k \cdot 2^k + k \cdot 2^k + |xs| & \\
= & (k + 1) \cdot 2^{k+1} & \text{by assumption } |xs| = 2^{k+1}
\end{array}
$$

□

Exercise 2.6. The definition of *msort* is inefficient in that it calls *length*, *take* and *drop* for each list. Instead we can split the list into two halves by traversing it only once and putting its elements alternately on two piles, for example *halve* [2, 3, 4] ([0], [1]) = ([4, 2, 0], [3, 1]). Define *halve* and *msort2*

$$\begin{aligned}&msort2\ [] = []\\ &msort2\ [x] = [x]\\ &msort2\ xs\\ &= (\textbf{let}\ (ys1, ys2) = halve\ xs\ ([], [])\ \textbf{in}\ merge\ (msort2\ ys1)\ (msort2\ ys2))\end{aligned}$$

and prove $mset\ (msort2\ xs) = mset\ xs$ and $sorted\ (msort2\ xs)$. (Hint for Isabelle users: The definition of *msort2* is tricky because its termination relies on suitable properties of *halve*.)

2.5 Bottom-Up Merge Sort

Bottom-up merge sort starts by turning the input $[x_1, \ldots, x_n]$ into the list $[[x_1], \ldots, [x_n]]$. Then it passes over this list of lists repeatedly, merging pairs of adjacent lists on every pass until at most one list is left.

```
merge_adj :: 'a list list ⇒ 'a list list
merge_adj [] = []
merge_adj [xs] = [xs]
merge_adj (xs # ys # zss) = merge xs ys # merge_adj zss

merge_all :: 'a list list ⇒ 'a list
merge_all [] = []
merge_all [xs] = xs
merge_all xss = merge_all (merge_adj xss)

msort_bu :: 'a list ⇒ 'a list
msort_bu xs = merge_all (map (λx. [x]) xs)
```

Termination of *merge_all* relies on the fact that *merge_adj* halves the length of the list (rounding up). Computation induction proves

$$|merge_adj\ xs| = (|xs| + 1)\ div\ 2 \tag{2.16}$$

2.5.1 Functional Correctness

We introduce the abbreviation *mset_mset* :: *'a list list* ⇒ *'a multiset*:

$$mset_mset\ xss \equiv \sum_{\#}\ (image_mset\ mset\ (mset\ xss))$$

These are the key properties of the functions involved:

$$mset_mset\ (merge_adj\ xss) = mset_mset\ xss$$
$$mset\ (merge_all\ xss) = mset_mset\ xss \tag{2.17}$$
$$mset\ (msort_bu\ xs) = mset\ xs$$
$$Ball\ (set\ xss)\ sorted \longrightarrow Ball\ (set\ (merge_adj\ xss))\ sorted$$
$$Ball\ (set\ xss)\ sorted \longrightarrow sorted\ (merge_all\ xss) \tag{2.18}$$
$$sorted\ (msort_bu\ xs)$$

The third and the last property prove functional correctness of *msort_bu*. The proof of each proposition may use the preceding proposition and the propositions (2.10) and (2.12). The propositions about *merge_adj* and *merge_all* are proved by computation inductions.

2.5.2 Running Time Analysis

Again, we count only comparisons:

$$C_{merge_adj} :: 'a\ list\ list \Rightarrow nat$$
$$C_{merge_adj}\ [] = 0$$
$$C_{merge_adj}\ [_] = 0$$
$$C_{merge_adj}\ (xs\ \#\ ys\ \#\ zss) = C_{merge}\ xs\ ys + C_{merge_adj}\ zss$$

$$C_{merge_all} :: 'a\ list\ list \Rightarrow nat$$
$$C_{merge_all}\ [] = 0$$
$$C_{merge_all}\ [xs] = 0$$
$$C_{merge_all}\ xss = C_{merge_adj}\ xss + C_{merge_all}\ (merge_adj\ xss)$$

$$C_{msort_bu} :: 'a\ list \Rightarrow nat$$
$$C_{msort_bu}\ xs = C_{merge_all}\ (map\ (\lambda x.\ [x])\ xs)$$

By simple computation inductions we obtain:

$$even\ |xss| \wedge (\forall\, xs \in set\ xss.\ |xs| = m) \longrightarrow$$
$$(\forall\, xs \in set\ (merge_adj\ xss).\ |xs| = 2 \cdot m) \tag{2.19}$$
$$(\forall\, xs \in set\ xss.\ |xs| = m) \longrightarrow C_{merge_adj}\ xss \leq m \cdot |xss| \tag{2.20}$$

using (2.13) for (2.19) and (2.15) for (2.20).

Lemma 2.5. $(\forall\, xs{\in}set\ xss.\ |xs| = m) \wedge |xss| = 2^k \longrightarrow$
$C_{merge_all}\ xss \leq m \cdot k \cdot 2^k$

Proof by induction on the computation of *merge_all*. We concentrate on the nontrivial recursive case arising from the third equation. We assume $|xss| > 1$, $\forall\, xs{\in}set\ xss.\ |xs| = m$ and $|xss| = 2^k$. Clearly $k \geq 1$ and thus $even\ |xss|$. Thus (2.19) implies $\forall\, xs{\in}set\ (merge_adj\ xss).\ |xs| = 2 \cdot m$. Also note

$$\begin{aligned}
&|merge_adj\ xss| \\
&= (|xss| + 1)\ div\ 2 && \text{using (2.16)} \\
&= 2^k - 1 && \text{using } |xss| = 2^k \text{ and } k \geq 1 \text{ by arithmetic}
\end{aligned}$$

Let $yss = merge_adj\ xss$. We can now prove the lemma:

$$\begin{aligned}
&C_{merge_all}\ xss = C_{merge_adj}\ xss + C_{merge_all}\ yss \\
&\leq m \cdot 2^k + C_{merge_all}\ yss && \text{using } |xss| = 2^k \text{ and (2.20)} \\
&\leq m \cdot 2^k + 2 \cdot m \cdot (k-1) \cdot 2^{k-1} \\
&\qquad \text{by IH using } \forall\, xs{\in}set\ yss.\ |xs| = 2 \cdot m \text{ and } |yss| = 2^{k-1} \\
&= m \cdot k \cdot 2^k && \square
\end{aligned}$$

Setting $m = 1$ we obtain the same upper bound as for top-down merge sort in Lemma 2.4:

Corollary 2.6. $|xs| = 2^k \longrightarrow C_{msort_bu}\ xs \leq k \cdot 2^k$

2.6 Natural Merge Sort

A disadvantage of all the sorting functions we have seen so far (except insertion sort) is that even in the best case they do not improve upon the $n \cdot \lg n$ bound. For example, given the sorted input [1, 2, 3, 4, 5], *msort_bu* will, as a first step, create [[1], [2], [3], [4], [5]] and then merge this list of lists recursively.

A slight variation of bottom-up merge sort, sometimes referred to as "natural merge sort," first partitions the input into its constituent ascending and descending subsequences (collectively referred to as **runs**) and only then starts merging. In the above example we would get *merge_all* [[1, 2, 3, 4, 5]], which returns immediately with the result [1, 2, 3, 4, 5]. Assuming that obtaining runs is of linear complexity, this yields a best-case performance that is linear in the number of list elements.

Function *runs* computes the initial list of lists; it is defined mutually recursively with *asc* and *desc*, which gather ascending and descending runs in accumulating parameters:

```
runs :: 'a list ⇒ 'a list list
runs (a # b # xs)
= (if b < a then desc b [a] xs else asc b ((#) a) xs)
runs [x] = [[x]]
runs [] = []

asc :: 'a ⇒ ('a list ⇒ 'a list) ⇒ 'a list ⇒ 'a list list
asc a as (b # bs)
= (if ¬ b < a then asc b (as ∘ (#) a) bs else as [a] # runs (b # bs))
asc a as [] = [as [a]]

desc :: 'a ⇒ 'a list ⇒ 'a list ⇒ 'a list list
desc a as (b # bs)
= (if b < a then desc b (a # as) bs else (a # as) # runs (b # bs))
desc a as [] = [a # as]
```

Function *desc* needs to reverse the descending run it collects. Therefore a natural choice for the type of its accumulator *as* is *list*, since recursively prepending elements (using (#)) ultimately yields a reversed list.

Function *asc* collects an ascending run and is slightly more complicated than *desc*. If we used lists, we could accumulate the elements similarly to *desc* but using *as* @ [*a*] instead of *a* # *as*. This would take quadratic time in the number of appended elements. Therefore the "standard" solution is to accumulate elements using (#) and to reverse the accumulator in linear time (as shown in Section 1.5.1) at the end. However, another interesting option (that yields better performance for some functional languages, like Haskell) is to use **difference lists**. This is the option we chose for *asc*.

In the functional programming world, difference lists are a well-known trick to append lists in constant time by representing lists as functions of type *'a list* ⇒ *'a list*. For difference lists, we have the following correspondences: empty list [] ≈ λ*x*. *x*, singleton list [*x*] ≈ (#) *x*, and list append *xs* @ *ys* ≈ *xs* ∘ *ys* (taking constant time). Moreover, transforming a difference list *xs* into a normal list is as easy as *xs* [] (taking linear time).

Note that, due to the mutually recursive definitions of *runs*, *asc*, and *desc*, whenever we prove a property of *runs*, we simultaneously have to prove suitable properties of *asc* and *desc* using mutual induction.

Natural merge sort is the composition of *merge_all* and *runs*:

```
nmsort :: 'a list ⇒ 'a list
nmsort xs = merge_all (runs xs)
```

2.6.1 Functional Correctness

We have

$$(\forall\, xs\ ys.\ f\ (xs\ @\ ys) = f\ xs\ @\ ys) \longrightarrow$$
$$mset_mset\ (asc\ x\ f\ ys) = \{\!\{x\}\!\} + mset\ (f\ []) + mset\ ys \tag{2.21}$$
$$mset_mset\ (desc\ x\ xs\ ys) = \{\!\{x\}\!\} + mset\ xs + mset\ ys \tag{2.22}$$
$$mset_mset\ (runs\ xs) = mset\ xs \tag{2.23}$$
$$mset\ (nmsort\ xs) = mset\ xs \tag{2.24}$$

where (2.23), (2.21), and (2.22) are proved simultaneously. The assumption of (2.21) on f ensures that f is a difference list. We use (2.23) together with (2.17) in order to show (2.24). Moreover, we have

$$Ball\ (set\ (runs\ xs))\ sorted \tag{2.25}$$
$$sorted\ (nmsort\ xs) \tag{2.26}$$

where we use (2.25) together with (2.18) to obtain (2.26).

2.6.2 Running Time Analysis

Once more, we only count comparisons:

```
C_runs :: 'a list ⇒ nat
C_runs (a # b # xs) = 1 + (if b < a then C_desc b xs else C_asc b xs)
C_runs [] = 0
C_runs [_] = 0

C_asc :: 'a ⇒ 'a list ⇒ nat
C_asc a (b # bs) = 1 + (if ¬ b < a then C_asc b bs else C_runs (b # bs))
C_asc _ [] = 0

C_desc :: 'a ⇒ 'a list ⇒ nat
C_desc a (b # bs) = 1 + (if b < a then C_desc b bs else C_runs (b # bs))
C_desc _ [] = 0
```

$$C_{nmsort} :: 'a\ list \Rightarrow nat$$
$$C_{nmsort}\ xs = C_{runs}\ xs + C_{merge_all}\ (runs\ xs)$$

Again note the mutually recursive definitions of C_{runs}, C_{asc}, and C_{desc}. Hence the remark on proofs about *runs* also applies to proofs about C_{runs}.

Before talking about C_{nmsort}, we need a variant of Lemma 2.5 that also works for lists whose lengths are not powers of two (since the result of *runs* will usually not satisfy this property).

To this end, we will need the following two results, which we prove by two simple computation inductions using (2.15) and (2.13):

$$C_{merge_adj}\ xss \leq |concat\ xss| \tag{2.27}$$
$$|concat\ (merge_adj\ xss)| = |concat\ xss| \tag{2.28}$$

Lemma 2.7. $C_{merge_all}\ xss \leq |concat\ xss| \cdot \lceil \lg\ |xss| \rceil$

Proof by induction on the computation of C_{merge_all}. We concentrate on the nontrivial recursive case arising from the third equation. It follows that xss is of the form $xs\ @\ ys\ @\ zss$. Further note that for all $n :: nat$:

$$2 \leq n \longrightarrow \lceil \lg\ n \rceil = \lceil \lg\ ((n - 1)\ div\ 2 + 1) \rceil + 1 \tag{2.29}$$

Now, let $m = |concat\ xss|$. Then we have

$$\begin{aligned}
&C_{merge_all}\ xss \\
&= C_{merge_adj}\ xss + C_{merge_all}\ (merge_adj\ xss) \\
&\leq m + C_{merge_all}\ (merge_adj\ xss) && \text{using (2.27)} \\
&\leq m + |concat\ (merge_adj\ xss)| \cdot \lceil \lg\ |merge_adj\ xss| \rceil && \text{by IH} \\
&= m + m \cdot \lceil \lg\ |merge_adj\ xss| \rceil && \text{by (2.28)} \\
&= m + m \cdot \lceil \lg\ ((|xss| + 1)\ div\ 2) \rceil && \text{by (2.16)} \\
&= m + m \cdot \lceil \lg\ ((|xss| + 1)\ div\ 2 + 1) \rceil \\
&= m \cdot (\lceil \lg\ ((|xss| + 1)\ div\ 2 + 1) \rceil + 1) \\
&= m \cdot \lceil \lg\ (|zss| + 2) \rceil && \text{by (2.29)} \\
&= m \cdot \lceil \lg\ |xss| \rceil && \square
\end{aligned}$$

Three simple computation inductions, each performed simultaneously for the corresponding mutually recursive definitions, yield:

$$\begin{aligned}
&(\forall\ xs\ ys.\ f\ (xs\ @\ ys) = f\ xs\ @\ ys) \longrightarrow \\
&|concat\ (asc\ a\ f\ ys)| = 1 + |f\ []| + |ys|, \\
&|concat\ (desc\ a\ xs\ ys)| = 1 + |xs| + |ys|, \\
&|concat\ (runs\ xs)| = |xs| \qquad (2.30) \\
&(\forall\ xs\ ys.\ f\ (xs\ @\ ys) = f\ xs\ @\ ys) \longrightarrow |asc\ a\ f\ ys| \leq 1 + |ys|,
\end{aligned}$$

$$|desc\ a\ xs\ ys| \le 1 + |ys|,\ |runs\ xs| \le |xs| \tag{2.31}$$

$$C_{asc}\ a\ ys \le |ys|,\ C_{desc}\ a\ ys \le |ys|,\ C_{runs}\ xs \le |xs| - 1 \tag{2.32}$$

At this point we obtain an upper bound on the number of comparisons required by C_{nmsort}.

Lemma 2.8. $|xs| = n \longrightarrow C_{nmsort}\ xs \le n + n \cdot \lceil \lg\ n \rceil$

Proof. Note that

$$C_{merge_all}\ (runs\ xs) \le n \cdot \lceil \lg\ n \rceil \tag{$\star$}$$

as shown by the derivation:

$C_{merge_all}\ (runs\ xs)$
$\le |concat\ (runs\ xs)| \cdot \lceil \lg\ |runs\ xs| \rceil$ by Lemma 2.7 with $xss = runs\ xs$
$\le n \cdot \lceil \lg\ |runs\ xs| \rceil$ by (2.30)
$\le n \cdot \lceil \lg\ n \rceil$ by (2.31)

We conclude the proof by:

$C_{nmsort}\ xs = C_{runs}\ xs + C_{merge_all}\ (runs\ xs)$
$\le n + n \cdot \lceil \lg\ n \rceil$ using (2.32) and ($\star$) □

2.7 Stability

A sorting function is called **stable** if the order of equal elements is preserved. However, this only makes a difference if elements are not identified with their keys, as we have done so far. Let us assume instead that sorting is parameterized with a key function $f :: 'a \Rightarrow 'k$ that maps an element to its key and that the keys $'k$ are linearly ordered, not the elements. This is the specification of a sorting function *sort_key*:

$$mset\ (sort_key\ f\ xs) = mset\ xs$$
$$sorted\ (map\ f\ (sort_key\ f\ xs))$$

Assuming (for simplicity) we are sorting pairs of keys and some attached information, stability means that sorting $[(2,\ x),\ (1,\ z),\ (1,\ y)]$ yields $[(1,\ z),\ (1,\ y),\ (2,\ x)]$ and not $[(1,\ y),\ (1,\ z),\ (2,\ x)]$. That is, if we extract all elements with the same key *after* sorting xs, they should be in the same order as in xs:

$$filter\ (\lambda y.\ f\ y = k)\ (sort_key\ f\ xs) = filter\ (\lambda y.\ f\ y = k)\ xs$$

We will now define insertion sort adapted to keys and verify its correctness and stability.

```
insort_key :: ('b ⇒ 'a) ⇒ 'b ⇒ 'b list ⇒ 'b list
insort_key _ x [] = [x]
insort_key f x (y # ys)
= (if f x ≤ f y then x # y # ys else y # insort_key f x ys)

isort_key :: ('a ⇒ 'k) ⇒ 'a list ⇒ 'a list
isort_key _ [] = []
isort_key f (x # xs) = insort_key f x (isort_key f xs)
```

The proofs of the functional correctness properties

$$mset\ (isort_key\ f\ xs) = mset\ xs$$
$$sorted\ (map\ f\ (isort_key\ f\ xs)) \tag{2.33}$$

are completely analogous to their counterparts for plain *isort*.

The proof of stability uses three auxiliary properties:

$$(\forall x{\in}set\ xs.\ f\ a \le f\ x) \longrightarrow insort_key\ f\ a\ xs = a\ \#\ xs \tag{2.34}$$
$$\neg\ P\ x \longrightarrow filter\ P\ (insort_key\ f\ x\ xs) = filter\ P\ xs \tag{2.35}$$
$$sorted\ (map\ f\ xs) \wedge P\ x \longrightarrow$$
$$filter\ P\ (insort_key\ f\ x\ xs) = insort_key\ f\ x\ (filter\ P\ xs) \tag{2.36}$$

The first one is proved by a case analysis on xs. The other two are proved by induction on xs, using (2.34) in the proof of (2.36).

Lemma 2.9 (Stability of *isort_key*).
$filter\ (\lambda y.\ f\ y = k)\ (isort_key\ f\ xs) = filter\ (\lambda y.\ f\ y = k)\ xs$

Proof. The proof is by induction on xs. The base case is trivial. In the induction step we consider the list $a\ \#\ xs$ and perform a case analysis. If $f\ a \neq k$ the claim follows by IH using (2.35). Now assume $f\ a = k$:

$$\begin{aligned}
&filter\ (\lambda y.\ f\ y = k)\ (isort_key\ f\ (a\ \#\ xs))\\
&= filter\ (\lambda y.\ f\ y = k)\ (insort_key\ f\ a\ (isort_key\ f\ xs))\\
&\qquad\qquad \text{using } f\ a = k, (2.36), (2.33)\\
&= insort_key\ f\ a\ (filter\ (\lambda y.\ f\ y = k)\ xs) \qquad \text{by IH}\\
&= a\ \#\ filter\ (\lambda y.\ f\ y = k)\ xs \qquad \text{using } f\ a = k \text{ and } (2.34)\\
&= filter\ (\lambda y.\ f\ y = k)\ (a\ \#\ xs) \qquad \text{using } f\ a = k
\end{aligned}$$

□

As exercises we recommend to adapt some of the other sorting algorithms above to sorting with keys and to prove their correctness and stability.

3

Selection

A topic that is somewhat related to that of sorting is *selection*: given a list xs of length n with some linear ordering defined on it and a natural number $k < n$, return the k-th smallest number in the list. If xs is sorted, this is exactly the k-th element of the list.

The defining properties of the selection operation are as follows:

$$\begin{aligned} k < |xs| \longrightarrow & |\{\!\{y \in_\# mset\ xs \mid y < select\ k\ xs\}\!\}| \leq k \\ \wedge\ & |\{\!\{y \in_\# mset\ xs \mid y > select\ k\ xs\}\!\}| < |xs| - k \end{aligned} \tag{3.1}$$

These properties uniquely define the selection property, as shown by the following theorem:

Theorem 3.1 (Uniqueness of the selection operation).
If $k < |xs|$ *and*

$$\begin{aligned} &|\{\!\{z \in_\# mset\ xs \mid z < x\}\!\}| \leq k & &|\{\!\{z \in_\# mset\ xs \mid z > x\}\!\}| < |xs| - k \\ &|\{\!\{z \in_\# mset\ xs \mid z < y\}\!\}| \leq k & &|\{\!\{z \in_\# mset\ xs \mid z > y\}\!\}| < |xs| - k \end{aligned}$$

then $x = y$.

Proof. Suppose $x \neq y$ and then w.l.o.g. $x < y$. This implies:

$$\{\!\{z \in_\# mset\ xs \mid z \leq x\}\!\} \subseteq_\# \{\!\{z \in_\# mset\ xs \mid z < y\}\!\} \tag{3.2}$$

From this we can prove the contradiction $|xs| < |xs|$:

$$\begin{aligned} |xs| &= |\{\!\{z \in_\# mset\ xs \mid z \leq x\}\!\}| + |\{\!\{z \in_\# mset\ xs \mid z > x\}\!\}| \\ &\leq |\{\!\{z \in_\# mset\ xs \mid z < y\}\!\}| + |\{\!\{z \in_\# mset\ xs \mid z > x\}\!\}| && \text{using (3.2)} \\ &< k + (|xs| - k) && \text{using (3.1)} \\ &= |xs| \end{aligned}$$

□

An equivalent, more concrete definition is the following:

$$\begin{aligned} &\mathit{select} :: \mathit{nat} \Rightarrow \mathit{'a\ list} \Rightarrow \mathit{'a} \\ &\mathit{select}\ k\ \mathit{xs} = \mathit{sort}\ \mathit{xs}\ !\ k \end{aligned} \tag{3.3}$$

Theorem 3.2. *Our definition of select satisfies the conditions* (3.1).

Proof. If ys is a sorted list, the following can be proved by a straightforward induction on ys:

$$\{\!\!\{x \in_\# \mathit{mset}\ ys \mid x < ys\ !\ k\}\!\!\} \subseteq_\# \mathit{mset}\ (\mathit{take}\ k\ ys)\{\!\!\{x \in_\# \mathit{mset}\ ys \mid x > ys\ !\ k\}\!\!\} \subseteq_\# \mathit{mset}\ (\mathit{drop}\ (k+1)\ ys)$$

Taking the size of these multisets on both sides, we obtain:

$$|\{\!\!\{x \in_\# \mathit{mset}\ ys \mid x < ys\ !\ k\}\!\!\}| \leq k |\{\!\!\{x \in_\# \mathit{mset}\ ys \mid x > ys\ !\ k\}\!\!\}| < |ys| - k$$

By setting $ys := \mathit{sort}\ ys$ here, we have

$$\begin{aligned} & k \geq |\{\!\!\{x \in_\# \mathit{mset}\ (\mathit{sort}\ \mathit{xs}) \mid x < \mathit{sort}\ \mathit{xs}\ !\ k\}\!\!\}| \\ &= |\{\!\!\{x \in_\# \mathit{mset}\ \mathit{xs} \mid x < \mathit{sort}\ \mathit{xs}\ !\ k\}\!\!\}| && \text{using } \mathit{mset}\ (\mathit{sort}\ \mathit{xs}) = \mathit{mset}\ \mathit{xs} \\ &= |\{\!\!\{x \in_\# \mathit{mset}\ \mathit{xs} \mid x < \mathit{select}\ k\ \mathit{xs}\}\!\!\}| && \text{using (3.3)} \end{aligned}$$

and analogously for the elements greater than $\mathit{select}\ k\ \mathit{xs}$. □

We will frequently need two more important lemmas about *sort* and *select*, namely that they are invariant under permutation of the input list:

$$\mathit{mset}\ \mathit{xs} = \mathit{mset}\ ys \longrightarrow \mathit{sort}\ \mathit{xs} = \mathit{sort}\ ys \tag{3.4}$$
$$\mathit{mset}\ \mathit{xs} = \mathit{mset}\ ys \longrightarrow \mathit{select}\ k\ \mathit{xs} = \mathit{select}\ k\ ys \tag{3.5}$$

Equation (3.4) follows directly from the uniqueness of the *sort* operation, and (3.5) then follows from (3.4) and our definition of *select*.

The definition using $\mathit{sort}\ \mathit{xs}\ !\ k$ already gives us a straightforward $O(n \lg n)$ algorithm for the selection operation: sort the list with one of our $O(n \lg n)$ sorting algorithms and then return the k-th element of the resulting sorted list. It is also fairly easy to come up with an algorithm that has running time $O(kn)$, i.e. that runs in linear-time in n for any fixed k (see Exercise 3.3).

In the remainder of this chapter, we will look at a selection algorithm due to Blum *et al.* [10] that achieves $O(n)$ running time *unconditionally* (i.e. for any $k < n$). Since a selection algorithm must, in general, inspect each element at least once (see Exercise 3.4), this running time is asymptotically optimal.

3.0.1 Exercises

Exercise 3.1. A simple special case of selection is *select* 0 *xs*, i.e. the minimum. Implement a linear-time function *select0* such that

$$xs \neq [] \longrightarrow \mathit{select0}\ xs = \mathit{select}\ 0\ xs$$

and prove this. This function should be tail-recursive and traverse the list exactly once. You need not prove the linear running time (it should be obvious).

Exercise 3.2. How can your *select0* algorithm be modified to obtain an analogous algorithm *select1* such that

$$|xs| > 1 \longrightarrow \mathit{select1}\ xs = \mathit{select}\ 1\ xs$$

Do not try to prove the correctness yet; it gets somewhat tedious and you will be able to prove it more easily after the next exercise.

Exercise 3.3.

1. Based on the previous two exercises, implement and prove correct an algorithm *select_fixed* that fulfils
   $$k < |xs| \longrightarrow \mathit{select_fixed}\ k\ xs = \mathit{select}\ k\ xs$$
   The algorithm must be tail-recursive with running time $O(kn)$ and traverse the list exactly once.
   Hint: one approach is to first define a function *take_sort* that computes $\mathit{take}\ m\ (\mathit{sort}\ xs)$ in time $O(mn)$.
2. Prove your *select1* from the previous exercise correct by showing that it is equivalent to *select_fixed* 1.
3. Define a suitable time function for your *select_fixed*. Prove that this time function is $O(kn)$, i.e. that
   $$T_{\mathit{select_fixed}}\ k\ xs \leq C_1 \cdot k \cdot |xs| + C_2 \cdot |xs| + C_3 \cdot k + C_4$$
   for all $k < |xs|$ for some constants C_1 to C_4.
   If you have trouble finding the concrete values for these constants, try proving the result with symbolic constants first and observe what conditions need to be fulfilled in order to make the induction step go through.

Exercise 3.4. Show that if *xs* is a list of integers with no repeated elements, an algorithm computing the result of $\mathit{select}\ k\ xs$ must examine every single element, i.e. for any index $i < |xs|$, the i-th element can be replaced by some other number such that the result changes. Formally:

$$k < |xs| \wedge i < |xs| \wedge \mathit{distinct}\ xs \longrightarrow$$
$$(\exists\, z.\ \mathit{distinct}\ (xs[i := z]) \wedge \mathit{select}\ k\ (xs[i := z]) \neq \mathit{select}\ k\ xs)$$

Here, the notation $xs[i := z]$ denotes the list *xs* where the i-th element has been replaced with z (the first list element, as always, having index 0).

Hint: a lemma you might find useful is that $\lambda k.\ \mathit{select}\ k\ xs$ is injective if *xs* has no repeated elements.

3.1 A Divide-and-Conquer Approach

As a first step in our attempt to derive an efficient algorithm for selection, recall what we did with the function *partition*3 in the threeway quicksort algorithm in Exercise 2.4: we picked some pivot x from xs and partitioned the input list xs into the sublists ls, es, and gs that were smaller, equal, and greater than x, respectively.

If we do the same for *select* k xs, there are three possible cases:

- If $k < |ls|$, the element we are looking for is located in ls. To be more precise, it is the k-th smallest element of ls, i.e. *select* k ls.
- If $k < |ls| + |es|$, the element we are looking for is located in es and must therefore be x itself.
- Otherwise, the element we are looking for must be located in gs. More precisely, it is the k'-th smallest element of gs where $k' = k - |ls| - |es|$.

This gives us a straightforward recursive divide-and-conquer algorithm for selection. To prove this correct, we first prove the following lemma about *select*:

$$\begin{aligned}&k < |ys| + |zs| \wedge (\forall\, y{\in}set\ ys.\ \forall\, z{\in}set\ zs.\ y \leq z) \longrightarrow\\ &select\ k\ (ys\ @\ zs)\\ &= (\textbf{if}\ k < |ys|\ \textbf{then}\ select\ k\ ys\ \textbf{else}\ select\ (k - |ys|)\ zs)\end{aligned} \tag{3.6}$$

Proof. The assumptions imply that $sort\ xs\ @\ sort\ ys$ is sorted, so that due to the uniqueness of the *sort* operation, we have:

$$sort\ (xs\ @\ ys) = sort\ xs\ @\ sort\ ys \tag{3.7}$$

Then:

$$\begin{aligned}&select\ k\ (xs\ @\ ys)\\ &= sort\ (xs\ @\ ys)\ !\ k && \text{using (3.3)}\\ &= (sort\ xs\ @\ sort\ ys)\ !\ k && \text{using (3.7)}\\ &= \textbf{if}\ k < |xs|\ \textbf{then}\ sort\ xs\ !\ k\ \textbf{else}\ sort\ ys\ !\ (k - |xs|)\\ &= \textbf{if}\ k < |xs|\ \textbf{then}\ select\ k\ xs\ \textbf{else}\ select\ (k - |xs|)\ ys && \text{using (3.3)}\end{aligned}$$

□

Now the recurrence outlined before is a direct consequence:

Theorem 3.3 (A recurrence for *select*).

$$\begin{aligned}&k < |xs| \longrightarrow\\ &select\ k\ xs\\ &= (\textbf{let}\ (ls,\ es,\ gs) = partition3\ x\ xs\\ &\quad \textbf{in if}\ k < |ls|\ \textbf{then}\ select\ k\ ls\\ &\qquad \textbf{else if}\ k < |ls| + |es|\ \textbf{then}\ x\ \textbf{else}\ select\ (k - |ls| - |es|)\ gs)\end{aligned}$$

Proof. We clearly have $mset\ xs = mset\ ls + mset\ es + mset\ gs$ and $|xs| = |ls| + |es| + |gs|$. Then:

$$\begin{aligned}
&select\ k\ xs \\
&= select\ k\ (ls\ @\ es\ @\ gs) && \text{using (3.5)} \\
&= \textbf{if}\ k < |ls|\ \textbf{then}\ select\ k\ ls \\
&\quad \textbf{else if}\ k - |ls| < |es|\ \textbf{then}\ select\ (k - |ls|)\ es && \text{using (3.6)} \\
&\quad \textbf{else}\ select\ (k - |ls| - |es|)\ gs
\end{aligned}$$

Clearly, $k - |ls| < |es| \longleftrightarrow k < |ls| + |es|$ and $select\ (k - |ls|)\ es = x$ since $select\ (k - |ls|)\ es \in set\ es$ and $set\ es = \{x\}$ by construction. □

Note that this holds for *any* pivot x. Indeed, x need not even be in the list itself. Therefore, the algorithm is (partially) correct no matter what pivot we choose. However, the number of recursive calls (and thereby the running time) depends strongly on the pivot choice:

- If we always choose a pivot that is smaller than any element in the list, the algorithm does not terminate at all.
- If we choose the smallest element in the list as a pivot every time, only one element is removed from the list in every recursion step so that we get n recursive calls in total. Since we do a linear amount of work in every step, this leads to a running time of $\Theta(n^2)$.
- If we choose pivots from the list at random, the worst-case running time is again $\Theta(n^2)$, but the expected running time can be shown to be $\Theta(n)$, similarly to the situation in Quicksort. Indeed, it can also be shown that it is very unlikely that the running time is significantly "worse than linear" [39, Section 2.5].
- If we choose a pivot that cuts the list in half every time, we get roughly $\lg n$ recursive steps and, by the Master Theorem, a running time of $\Theta(n)$ (assuming we can find such a pivot in linear time).

Clearly, the last case is the most desirable one. The element that cuts the list in half is called the *median* (a term widely used in statistics, albeit with a slightly different meaning). For lists of odd length, there is a unique element that achieves this, whereas for lists of even length there are two such elements (e.g. for the list [1,2,3,4], both 2 and 3 work). For our purposes, we always use the smaller one and define:

```
median :: 'a list ⇒ 'a
median xs = select ((|xs| - 1) div 2) xs
```

Unfortunately, computing the median of a list is no easier than selection (see Exercise 3.5), so it seems that, for now, this does not really help us.

Exercise 3.5. Show that *select k xs* can be reduced to computing the median of a list in linear time, i.e. give a linear-time function *reduce_select_median* such that

$$xs \neq [] \wedge k < |xs| \longrightarrow$$
$$reduce_select_median\ k\ xs \neq [] \wedge$$
$$median\ (reduce_select_median\ k\ xs) = select\ k\ xs$$

3.2 The Median of Medians

Since, as we have seen, computing a true median in every recursive step is too expensive, the next natural question is: is there something that is *almost* as good as a median but easier to compute?

This is indeed the case, and this is where the ingenuity of the algorithm is hidden: instead of computing the median of *all* the list elements, we use another list element M, which is computed as follows:

- chop the list into groups of 5 elements each (possibly with one smaller group at the end if n is not a multiple of 5)
- compute the median of each of the $\lceil \frac{n}{5} \rceil$ groups (which can be done in constant time for each group using e.g. insertion sort, since their sizes are bounded by 5)
- compute the median M of these $\lceil \frac{n}{5} \rceil$ elements (which can be done by a recursive call to the selection algorithm)

We call M the *median of medians*. M is not quite as good a pivot as the true median, but it is still fairly decent:

Theorem 3.4 (Pivoting bounds for the median of medians). *Let xs be a list and let $\prec$ be either $<$ or $>$. Define*

$$M := median\ (map\ median\ (chop\ 5\ xs))\ .$$

where the chop function performs the chopping described earlier. Then:

$$|\{\!\{y \in_\# mset\ xs \mid y \prec M\}\!\}| \leq \lceil 0.7 \cdot n + 3 \rceil$$

Proof. The result of *chop* 5 *xs* is a list of $\lceil n\ /\ 5 \rceil$ chunks, each of size at most 5, i.e. $|chop\ 5\ xs| = \lceil n\ /\ 5 \rceil$. Let us split these chunks into two groups according to whether their median is $\prec M$ or $\succeq M$:

$$Y_\prec := \{\!\{ys \in_\# mset\ (chop\ 5\ xs) \mid median\ ys \prec M\}\!\}$$
$$Y_\succeq := \{\!\{ys \in_\# mset\ (chop\ 5\ xs) \mid median\ ys \succeq M\}\!\}$$

We clearly have

$$mset\ xs = (\sum_{ys \leftarrow chop\ 5\ xs} mset\ ys) \tag{3.8}$$
$$mset\ (chop\ 5\ xs) = Y_{\prec} + Y_{\succeq} \tag{3.9}$$
$$\lceil n\ /\ 5 \rceil = |Y_{\prec}| + |Y_{\succeq}| \tag{3.10}$$

and since M is the median of the medians of the groups, we also know that:

$$|Y_{\prec}| < \tfrac{1}{2} \cdot \lceil n\ /\ 5 \rceil \tag{3.11}$$

The core idea of the proof is that any group $ys \in_{\#} Y_{\succeq}$ can have at most 2 elements that are $\prec M$:

$$\begin{aligned}
&|\{\!|y \in_{\#} mset\ ys \mid y \prec M|\!\}| \\
&\leq |\{\!|y \in_{\#} mset\ ys \mid y \prec median\ ys|\!\}| && \text{because } median\ ys \succeq M \\
&\leq |ys|\ div\ 2 && \text{using (3.1)} \\
&\leq 5\ div\ 2 = 2
\end{aligned}$$

And of course, since each group has size at most 5, any group in $ys \in_{\#} Y_{\prec}$ can contribute at most 5 elements. In summary, we have:

$$\begin{aligned}
&\forall ys \in_{\#} Y_{\prec}.\ |\{\!|y \in_{\#} mset\ ys \mid y \prec M|\!\}| \leq 2 \\
&\forall ys \in_{\#} Y_{\succeq}.\ |\{\!|y \in_{\#} mset\ ys \mid y \prec M|\!\}| \leq 5
\end{aligned} \tag{3.12}$$

With this, we can begin our estimation of the number of elements $\prec M$:

$$\begin{aligned}
&\{\!|y \in_{\#} mset\ xs \mid y \prec M|\!\} \\
&= \{\!|y \in_{\#} (\sum_{ys \leftarrow chop\ 5\ xs} mset\ ys) \mid y \prec M|\!\} && \text{using(3.8)} \\
&= \sum_{ys \leftarrow chop\ 5\ xs} \{\!|y \in_{\#} mset\ ys \mid y \prec M|\!\} \\
&= \sum_{ys \in_{\#} (Y_{\prec} + Y_{\succeq})} \{\!|y \in_{\#} mset\ ys \mid y \prec M|\!\} && \text{using (3.9)}
\end{aligned}$$

Taking the size of both sides, we have

$$\begin{aligned}
&|\{\!|y \in_{\#} mset\ xs \mid y \prec M|\!\}| \\
&\leq \sum_{ys \in_{\#} (Y_{\prec} + Y_{\succeq})} |\{\!|y \in_{\#} mset\ ys \mid y \prec M|\!\}| \\
&= \sum_{ys \in_{\#} Y_{\prec}} |\{\!|y \in_{\#} mset\ ys \mid y \prec M|\!\}| + \\
&\quad \sum_{ys \in_{\#} Y_{\succeq}} |\{\!|y \in_{\#} mset\ ys \mid y \prec M|\!\}| \\
&\leq (\sum_{ys \in_{\#} Y_{\prec}} 5) + (\sum_{ys \in_{\#} Y_{\succeq}} 2) && \text{using (3.12)} \\
&= 5 \cdot |Y_{\prec}| + 2 \cdot |Y_{\succeq}| \\
&= 2 \cdot \lceil n\ /\ 5 \rceil + 3 \cdot |Y_{\succeq}| && \text{using (3.10)} \\
&\leq 3.5 \cdot \lceil n\ /\ 5 \rceil && \text{using (3.11)} \\
&= 3.5 \cdot \lceil n\ /\ 5 \rceil \leq \lceil 0.7 \cdot n + 3 \rceil
\end{aligned}$$

The delicate arithmetic reasoning about rounding in the end can thankfully be done fully automatically by Isabelle's `linarith` tactic. □

3.3 Selection in Linear Time

We now have all the ingredients to write down our algorithm: the base cases (i.e. sufficiently short lists) can be handled using the naive approach of performing insertion sort and then returning the k-th element. For bigger lists, we perform the divide-and-conquer approach outlined in Theorem 3.3 using M as a pivot. We have two recursive calls: one on a list with exactly $\lceil 0.2 \cdot n \rceil$ elements to compute M, and one on a list with at most $\lceil 0.7 \cdot n + 3 \rceil$. We will still need to show later that this actually leads to a linear-time algorithm, but the fact that $0.7 + 0.2 < 1$ is at least encouraging.

The full algorithm looks like this:

```
chop :: nat ⇒ 'a list ⇒ 'a list list
chop 0 _ = []
chop _ [] = []
chop n xs = take n xs # chop n (drop n xs)

slow_select :: nat ⇒ 'a list ⇒ 'a
slow_select k xs = isort xs ! k

slow_median :: 'a list ⇒ 'a
slow_median xs = slow_select ((|xs| − 1) div 2) xs

mom_select :: nat ⇒ 'a list ⇒ 'a
mom_select k xs
= (if |xs| ≤ 20 then slow_select k xs
   else let M = mom_select ((⌈|xs| / 5⌉ − 1) div 2)
                  (map slow_median (chop 5 xs));
            (ls, es, gs) = partition3 M xs
        in if k < |ls| then mom_select k ls
           else if k < |ls| + |es| then M
           else mom_select (k − |ls| − |es|) gs)
```

The (partial) correctness statement is simple:

Theorem 3.5 (Partial Correctness of *mom_select*). *Let xs be a list and $k < |xs|$. Then if mom_select k xs terminates, we have*

$$mom_select\ k\ xs = select\ k\ xs\ .$$

Proof. Straightforward computation induction using Theorem 3.3. □

Theorem 3.6 (Termination of *mom_select*). *Let xs be a list and $k < |xs|$. Then mom_select k xs terminates.*

Proof. We prove termination by using the length of the list xs as a termination measure. We perform a computation induction and need to show that this length decreases in every recursive call if $20 < |xs|$. This is easy to see:

- The list in the first recursive call has length $\lceil |xs| \;/\; 5 \rceil$, which is strictly less than $|xs|$ if $1 < |xs|$.
- The length of the list in the second recursive call is at most $|xs| - 1$: by induction hypothesis, the first recursive call terminates, so by Theorem 3.5 we know that $M = median\ (map\ median\ (chop\ 5\ xs))$ and thus:
  $M \in set\ (map\ median\ (chop\ 5\ xs))$
  $= \{median\ ys \mid ys \in set\ (chop\ 5\ xs)\}$
  $\subseteq \bigcup_{ys \in set\ (chop\ 5\ xs)} set\ ys = set\ xs$
  Hence, $M \in set\ xs$ but $M \notin set\ ls$ and $M \notin set\ gs$ by construction.

□

Exercise 3.6. The recursive definition of *mom_select* handles the cases $|xs| \leq 20$ through the naive algorithm using insertion sort. The constant 20 here seems somewhat arbitrary. Find the smallest constant n_0 for which the algorithm still works. Why do you think 20 was chosen?

3.4 Time Functions

It remains to show now that this indeed leads to a linear-time algorithm. The time function for our selection algorithm is as follows:

```
T_mom_select :: nat ⇒ 'a list ⇒ nat
T_mom_select k xs
= (if |xs| ≤ 20 then T_slow_select k xs
   else let xss = chop 5 xs;
            ms = map slow_median xss;
            idx = ((|xs| + 4) div 5 − 1) div 2;
            x = mom_select idx ms;
            (ls, es, gs) = partition3 x xs
        in T_mom_select idx ms + T_chop 5 xs + T_map T_slow_median xss +
           T_partition3 x xs + T_length ls + T_length es + 1 +
           (if k < |ls| then T_mom_select k ls
            else if k < |ls| + |es| then 0
            else T_mom_select (k − |ls| − |es|) gs))
```

We can then prove

$$T_{mom_select}\ k\ xs \leq T'_{mom_select}\ |xs|$$

where the upper bound T'_{mom_select} is defined as follows:

$$\begin{aligned}
&T'_{mom_select} :: nat \Rightarrow nat\\
&T'_{mom_select}\ n\\
&= (\textbf{if } n \leq 20 \textbf{ then } 463\\
&\quad \textbf{else } T'_{mom_select}\ \lceil 0.2 \cdot n \rceil + T'_{mom_select}\ \lceil 0.7 \cdot n + 3 \rceil + 17 \cdot n + 50)
\end{aligned}$$

The time functions of the auxiliary functions used here can be found in Section B.2 in the appendix. The proof is a simple computation induction using Theorem 3.4 and the time bounds for the auxiliary functions from Chapter B in the appendix.

The next section will be dedicated to showing that $T'_{mom_select} \in O(n)$.

Exercise 3.7. Show that the upper bound $\lceil 0.7 \cdot n + 3 \rceil$ is fairly tight by giving an infinite family $(xs_i)_{i \in \mathbb{N}}$ of lists with increasing lengths for which more than 70 % of the elements are larger than the median of medians (with chopping size 5). In Isabelle terms: define a function $f :: nat \Rightarrow nat\ list$ such that $\forall n.\ |f\ n| < |f\ (n+1)|$ and

$$\frac{|\{\!\{y \in_\# mset\ (f\ n) \mid y > mom\ (f\ n)\}\!\}|}{|f\ n|} > 0.7$$

where $mom\ xs = median\ (map\ median\ (chop\ 5\ xs))$.

3.5 Akra–Bazzi Light

The function T'_{mom_select} (let us write it as f for now) fulfils the recurrence

$$n > 20 \longrightarrow f\ n = f\ \lceil 0.2 \cdot n \rceil + f\ \lceil 0.7 \cdot n + 3 \rceil + 17 \cdot n + 50$$

Such divide-and-conquer recurrences are beyond the "normal" Master Theorem, but a generalisation, the Akra–Bazzi Theorem [3, 48, 17], does apply to them. Let us first abstract the situation a bit and consider the recurrence

$$n > 20 \longrightarrow f\ n = f\ \lceil a \cdot n + b \rceil + f\ \lceil c \cdot n + d \rceil + C_1 \cdot n + C_2$$

where $0 < a,\ b < 1$ and $C_1,\ C_2 > 0$. The Akra–Bazzi Theorem then tells us that such a function is $O(n)$ if (and only if) $a + b < 1$. We will prove one direction of this (the one we need) now – Akra–Bazzi Light, so to say.

Instead of presenting the full theorem statement and its proof right away, let us take a more explorative approach. What we want to prove in the end is that there are real constants $C_3 > 0$ and C_4 such that $f\ n \leq C_3 \cdot n + C_4$ for all n. Suppose these constants are fixed already and now we want to prove that the inequality holds for them. For simplicity of the presentation, we assume $b,\ d \geq 0$, but note that these assumptions are unnecessary and the proof still works for negative b and d if we replace b and d with $max\ 0\ b$ and $max\ 0\ d$.

The obvious approach to show this is by induction on n, following the structure of the recurrence above. To do this, we use *complete induction* (i.e. the induction hypothesis holds for all $n' < n$)[1] and a case analysis on $n > n_0$ (where n_0 is some constant we will determine later).

The two cases we have to show in the induction are then:

Base case: $\forall\, n{\leq}n_0.\ f\ n \leq C_3 \cdot n + C_4$
Step: $\forall\, n{>}n_0.\ (\forall\, m{<}n.\ f\ m \leq C_3 \cdot m + C_4) \longrightarrow f\ n \leq C_3 \cdot n + C_4$

We can see that in order to even be able to apply the induction hypothesis in the induction step, we need $\lceil a \cdot n + b \rceil < n$. We can make the estimate[2]

$$\lceil a \cdot n + b \rceil \leq a \cdot n + b + 1 \overset{!}{<} n$$

and then solve for n, which gives us $n \overset{!}{>} \frac{b+1}{1-a}$. If we do the same for c and d as well, we get the conditions

$$n_0 \geq \frac{b+1}{1-a} \quad \text{and} \quad n_0 \geq \frac{d+1}{1-c} \tag{3.13}$$

However, it will later turn out that these are implied by the other conditions we will have accumulated anyway.

Now that we have ensured that the basic structure of our induction will work out, let us continue with the two cases.

The base cases ($n \leq n_0$) is fairly uninteresting: we can simply choose C_4 to be large enough to satisfy all of them.

In the recursive step, unfolding one step of the recurrence and applying the induction hypotheses leaves us with the following proof obligation

$$C_3 \cdot \lceil a \cdot n + b \rceil + C_4 + (C_3 \cdot \lceil c \cdot n + d \rceil + C_4) + C_1 \cdot n + C_2 \overset{!}{\leq} C_3 \cdot n + C_4\ ,$$

[1] In Isabelle, the corresponding rule is called `less_induct`.

[2] The notation $\overset{!}{<}$ stands for "must be less than". It emphasises that this inequality is not a consequence of what we have shown so far, but something that we still need to show, or in this case something that we need to ensure by adding suitable preconditions.

or, equivalently,

$$C_3 \cdot (\lceil a \cdot n + b \rceil + \lceil c \cdot n + d \rceil - n) + C_1 \cdot n + C_2 + C_4 \overset{!}{\leq} 0 \ ,$$

We estimate the left-hand side like this:

$$\begin{aligned}
&C_3 \cdot (\lceil a \cdot n + b \rceil + \lceil c \cdot n + d \rceil - n) + C_1 \cdot n + C_2 + C_4 \\
&\leq C_3 \cdot (a \cdot n + b + 1 + (c \cdot n + d + 1) - n) + C_1 \cdot n + C_2 + C_4 \\
&= C_3 \cdot (b + d + 2) + C_2 + C_4 - (C_3 \cdot (1 - a - c) - C_1) \cdot n \quad (*) \\
&\leq C_3 \cdot (b + d + 2) + C_2 + C_4 - (C_3 \cdot (1 - a - c) - C_1) \cdot n_0 \quad (\dagger) \\
&\overset{!}{\leq} 0
\end{aligned}$$

The step from $(*)$ to $(\dagger)$ uses the fact that $n > n_0$ and requires the factor $C_3 \cdot (1 - a - c) - C_1$ in front of the n to be positive, i.e. we need to add the assumption

$$C_3 > \frac{C_1}{1 - a - c} \ . \tag{3.14}$$

The term $(\dagger)$ (which we want to be ≤ 0) is now a constant. If we solve that inequality for C_3, we get the following two additional conditions:

$$n_0 > \frac{b + d + 2}{1 - a - c} \quad \text{and} \quad C_3 \geq \frac{C_1 \cdot n_0 + C_2 + C_4}{(1 - a - c) \cdot n_0 - b - d - 2} \tag{3.15}$$

The former of these directly implies our earlier conditions (3.13), so we can safely discard those now.

Now all we have to do is to find a combinations of n_0, C_3, and C_4 that satisfy (3.14) and (3.15). This is straightforward:

$$n_1 := max\ n_0 \left(\left\lceil \frac{b + d + 2}{1 - a - c} \right\rceil + 1 \right) \qquad C_4 := Max\ \{f\ n \mid n \leq 20\}$$

$$C_3 := max \left(\frac{C_1}{1 - a - c} \right) \left(\frac{C_1 \cdot n_1 + C_2 + C_4}{(1 - a - c) \cdot n_1 - b - d - 2} \right)$$

And with that, the induction goes through and we get the following theorem:

Theorem 3.7 (Akra Bazzi Light).

$$\begin{aligned}
&a > 0 \wedge c > 0 \wedge a + c < 1 \wedge C_1 \geq 0 \wedge \\
&(\forall\, n{>}n_0.\ f\ n = f\ \lceil a \cdot n + b \rceil + f\ \lceil c \cdot n + d \rceil + C_1 \cdot n + C_2) \longrightarrow \\
&(\exists\, C_3\ C_4.\ \forall\, n.\ f\ n \leq C_3 \cdot n + C_4)
\end{aligned}$$

Exercise 3.8.

1. Suppose that instead of groups of 5, we now chop into groups of size $l \geq 1$. Prove a corresponding generalisation of Theorem 3.4.
2. Examine (on paper only): how does this affect correctness and running time of our selection algorithm? Why do you think $l = 5$ was chosen?

Part II

Search Trees

4

Binary Trees

Binary trees are defined as a recursive data type:

datatype $'a\ tree = Leaf \mid Node\ ('a\ tree)\ 'a\ ('a\ tree)$

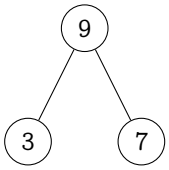
The following syntactic sugar is sprinkled on top:

$$\begin{aligned} \langle\rangle &\equiv Leaf \\ \langle l, x, r\rangle &\equiv Node\ l\ x\ r \end{aligned}$$

The trees l and r are the left and right **children** of the node $\langle l, x, r\rangle$.

Because most of our tree will be binary trees, we drop the "binary" most of the time and have also called the type merely *tree*.

When displaying a tree in the usual graphical manner we show only the *Node*s. For example, $\langle\langle\rangle, 3, \langle\rangle\rangle, 9, \langle\langle\rangle, 7, \langle\rangle\rangle\rangle$ is displayed like this:



The (label of the) **root** node is 9. The **depth** (or **level**) of some node (or leaf) in a tree is the distance from the root. We use these concepts only informally.

4.1 Basic Functions

Two canonical functions on data types are *set* and *map*:

$$
\begin{aligned}
&set_tree :: 'a\ tree \Rightarrow 'a\ set\\
&set_tree\ \langle\rangle = \{\}\\
&set_tree\ \langle l,\ x,\ r\rangle = set_tree\ l \cup \{x\} \cup set_tree\ r\\
\\
&map_tree :: ('a \Rightarrow 'b) \Rightarrow 'a\ tree \Rightarrow 'b\ tree\\
&map_tree\ f\ \langle\rangle = \langle\rangle\\
&map_tree\ f\ \langle l,\ x,\ r\rangle = \langle map_tree\ f\ l,\ f\ x,\ map_tree\ f\ r\rangle
\end{aligned}
$$

The *inorder*, *preorder* and *postorder* traversals (we omit the latter) list the elements in a tree in a particular order:

$$
\begin{aligned}
&inorder :: 'a\ tree \Rightarrow 'a\ list\\
&inorder\ \langle\rangle = []\\
&inorder\ \langle l,\ x,\ r\rangle = inorder\ l\ @\ [x]\ @\ inorder\ r\\
\\
&preorder :: 'a\ tree \Rightarrow 'a\ list\\
&preorder\ \langle\rangle = []\\
&preorder\ \langle l,\ x,\ r\rangle = x\ \#\ preorder\ l\ @\ preorder\ r
\end{aligned}
$$

These two size functions count the number of *Node*s and *Leaf*s in a tree:

$$
\begin{aligned}
&size :: 'a\ tree \Rightarrow nat\\
&|\langle\rangle| = 0\\
&|\langle l,\ _,\ r\rangle| = |l| + |r| + 1\\
\\
&size1 :: 'a\ tree \Rightarrow nat\\
&|\langle\rangle|_1 = 1\\
&|\langle l,\ _,\ r\rangle|_1 = |l|_1 + |r|_1
\end{aligned}
$$

The syntactic sugar $|t|$ for $size\ t$ and $|t|_1$ for $size1\ t$ is only used in this text, not in the Isabelle theories.

Induction proves a convenient fact that explains the name *size*1:

$$|t|_1 = |t| + 1$$

The height and the minimal height of a tree are defined as follows:

```
height :: 'a tree ⇒ nat
h ⟨⟩ = 0
h ⟨l, _, r⟩ = max (h l) (h r) + 1

min_height :: 'a tree ⇒ nat
mh ⟨⟩ = 0
mh ⟨l, _, r⟩ = min (mh l) (mh r) + 1
```

You can think of them as the longest and shortest (cycle-free) path from the root to a leaf. The real names of these functions are *height* and *min_height*. The abbreviations h and mh are only used in this text, not in the Isabelle theories.

The obvious properties $h\ t \le |t|$ and $mh\ t \le h\ t$ and the following classical properties have easy inductive proofs:

$$2^{mh\ t} \le |t|_1 \qquad |t|_1 \le 2^{h\ t}$$

We will simply use these fundamental properties without referring to them by a name or number.

The set of subtrees of a tree is defined as follows:

```
subtrees :: 'a tree ⇒ 'a tree set
subtrees ⟨⟩ = {⟨⟩}
subtrees ⟨l, a, r⟩ = {⟨l, a, r⟩} ∪ subtrees l ∪ subtrees r
```

Note that every tree is a subtree of itself.

Exercise 4.1. Function *inorder* has quadratic complexity because the running time of (@) is linear in the length of its first argument. Define a function *inorder2* :: *'a tree* ⇒ *'a list* ⇒ *'a list* that avoids (@) but accumulates the result in its second parameter via (#) only. Its running time should be linear in the size of the tree. Prove *inorder2 t xs* = *inorder t* @ *xs*.

4.2 Complete Trees

A **complete tree** is one where all the leaves are on the same level. An example is shown in Figure 4.1. The predicate *complete* is defined recursively:

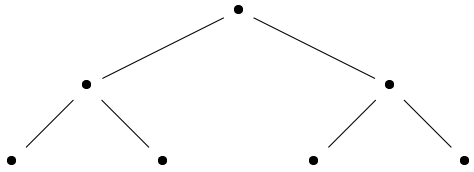


Fig. 4.1. A complete tree

```
complete :: 'a tree ⇒ bool
complete ⟨⟩ = True
complete ⟨l, _, r⟩ = (h l = h r ∧ complete l ∧ complete r)
```

This recursive definition is equivalent with the above definition that all leaves must have the same distance from the root. Formally:

Lemma 4.1. *complete* $t \longleftrightarrow mh\ t = h\ t$

Proof. The proof is by induction and case analyses on *min* and *max*. □

The following classic property of complete trees is easily proved by induction:

Lemma 4.2. *complete* $t \longrightarrow |t|_1 = 2^{h\ t}$

It turns out below that this is in fact a defining property of complete trees.

For complete trees we have $(2::'b)^{mh\ t} \le |t|_1 = 2^{h\ t}$. For incomplete trees the equalities the $=$ sign becomes a $<$ as the following two lemmas prove:

Lemma 4.3. $\neg$ *complete* $t \longrightarrow |t|_1 < 2^{h\ t}$

Proof. The proof is by induction. We focus in the induction step where $t = \langle l, x, r\rangle$. If t is incomplete, there are a number of cases and we prove $|t|_1 < 2^{h\ t}$ in each case. If $h\ l \neq h\ r$, consider the case $h\ l < h\ r$ (the case $h\ r < h\ l$ is symmetric). From $2^{h\ l} < 2^{h\ r}$, $|l|_1 \le 2^{h\ l}$ and $|r|_1 \le 2^{h\ r}$ the claim follows: $|t|_1 = |l|_1 + |r|_1 \le 2^{h\ l} + 2^{h\ r} < 2 \cdot 2^{h\ r} = 2^{h\ t}$. If $h\ l = h\ r$, then either l or r must be incomplete. We consider the case $\neg$ *complete* l (the case $\neg$ *complete* r is symmetric). From the IH $|l|_1 < 2^{h\ l}$, $|r|_1 \le 2^{h\ r}$ and $h\ l = h\ r$ the claim follows: $|t|_1 = |l|_1 + |r|_1 < 2^{h\ l} + 2^{h\ r} = 2 \cdot 2^{h\ r} = 2^{h\ t}$. □

Lemma 4.4. $\neg$ *complete* $t \longrightarrow 2^{mh\ t} < |t|_1$

The proof of this lemma is completely analogous to the previous proof except that one also needs to use Lemma 4.1.

From the contrapositive of Lemma 4.3 one obtains $|t|_1 = 2^{h\ t} \longrightarrow$ *complete* t, the converse of Lemma 4.2. Thus we arrive at:

Corollary 4.5. $complete\ t \longleftrightarrow |t|_1 = 2^{h\ t}$

The complete trees are precisely the ones where the height is exactly the logarithm of the number of leaves.

Exercise 4.2. Define a function *mcs* that computes a maximal complete subtree of some given tree. You are allowed only one traversal of the input but you may freely compute the height of trees and may even compare trees for equality. You are not allowed to use *complete* or *subtrees*.

Prove that *mcs* returns a complete subtree (which should be easy) and that it is maximal in height:

$$u \in subtrees\ t \wedge complete\ u \longrightarrow h\ u \leq h\ (mcs\ t)$$

Bonus: get rid of any tree equality tests in *mcs*.

4.3 Almost Complete Trees

An **almost complete tree** is one where the leaves may occur not just at the lowest level but also one level above:

```
acomplete :: 'a tree ⇒ bool
acomplete t = (h t − mh t ≤ 1)
```

An example of an almost complete tree is shown in Figure 4.2. You can think of an almost complete tree as a complete tree with (possibly) some additional nodes one level below the last full level.

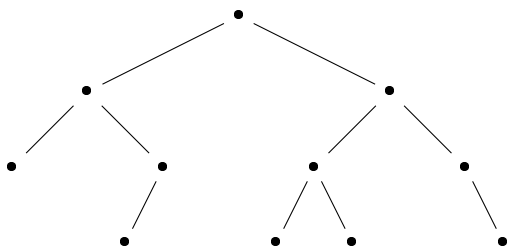


Fig. 4.2. Almost complete tree

Almost complete trees are important because among all the trees with the same number of nodes they have minimal height:

Lemma 4.6. $acomplete\ s \wedge |s| \leq |t| \longrightarrow h\ s \leq h\ t$

Proof. The proof is by cases. If *complete* s then, by Lemma 4.2, $2^{h\ s} = |s|_1 \leq |t|_1 \leq 2^{h\ t}$ and thus $h\ s \leq h\ t$. Now assume $\neg$ *complete* s. Then Lemma 4.4 yields $2^{mh\ s} < |s|_1 \leq |t|_1 \leq 2^{h\ t}$ and thus $mh\ s < h\ t$. Furthermore we have $h\ s - mh\ s \leq 1$ (from *acomplete* s), $h\ s \neq mh\ s$ (from Lemma 4.1) and $mh\ s \leq h\ s$, which together imply $mh\ s + 1 = h\ s$. With $mh\ s < h\ t$ this implies $h\ s \leq h\ t$. □

This is relevant for search trees because their height determines the worst case running time. Almost complete trees are optimal in that sense.

The following lemma yields an explicit formula for the height of an almost complete tree:

Lemma 4.7. $acomplete\ t \longrightarrow h\ t = \lceil \lg\ |t|_1 \rceil$

Proof. The proof is by cases. If t is complete, the claim follows from Lemma 4.2. Now assume t is incomplete. Then $h\ t = mh\ t + 1$ because *acomplete* t, $mh\ t \leq h\ t$ and *complete* $t \longleftrightarrow mh\ t = h\ t$ (Lemma 4.1). Together with $|t|_1 \leq 2^{h\ t}$ this yields $|t|_1 \leq 2^{mh\ t\ +\ 1}$ and thus $\lg\ |t|_1 \leq mh\ t + 1$. By Lemma 4.4 we obtain $mh\ t < \lg\ |t|_1$. These two bounds for $\lg\ |t|_1$ together imply the claimed $h\ t = \lceil \lg\ |t|_1 \rceil$. □

In the same manner we also obtain:

Lemma 4.8. $acomplete\ t \longrightarrow mh\ t = \lfloor \lg\ |t|_1 \rfloor$

4.3.1 Converting a List into an Almost Complete Tree

We will now see how to convert a list xs into an almost complete tree t such that $inorder\ t = xs$. If the list is sorted, the result is an almost complete binary search tree (see the next chapter). The basic idea is to cut the list in two halves, turn them into almost complete trees recursively and combine them. But cutting up the list in two halves explicitly would lead to an $n \cdot \lg\ n$ algorithm but we want a linear one. Therefore we use an additional *nat* parameter to tell us how much of the input list should be turned into a tree. The remaining list is returned with the tree:

```
bal :: nat ⇒ 'a list ⇒ 'a tree × 'a list
bal n xs
= (if n = 0 then (⟨⟩, xs)
    else let m = n div 2;
             (l, ys) = bal m xs;
```

```
        (r, zs) = bal (n − 1 − m) (tl ys)
    in (⟨l, hd ys, r⟩, zs))
```

The trick is not to chop xs in half but n because we assume that arithmetic is constant-time. Hence *bal* runs in linear time (see Exercise 4.4). Figure 4.3 shows the result of *bal* 10 [0..9].

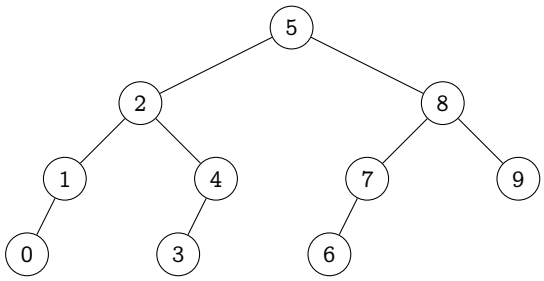


Fig. 4.3. Balancing [0, 1, 2, 3, 4, 5, 6, 7, 8, 9]

Balancing some prefix or all of a list or tree is easily derived:

```
bal_list :: nat ⇒ 'a list ⇒ 'a tree
bal_list n xs = fst (bal n xs)

balance_list :: 'a list ⇒ 'a tree
balance_list xs = bal_list |xs| xs

bal_tree :: nat ⇒ 'a tree ⇒ 'a tree
bal_tree n t = bal_list n (inorder t)

balance_tree :: 'a tree ⇒ 'a tree
balance_tree t = bal_tree |t| t
```

Correctness

The following lemma clearly expresses that *bal* n xs turns the prefix of length n of xs into a tree and returns the corresponding suffix of xs:

Lemma 4.9. $n \leq |xs| \wedge bal\ n\ xs = (t, zs) \longrightarrow$
$xs = inorder\ t\ @\ zs \wedge |t| = n$

Proof by complete induction on n, assuming that the property holds for all values below n. If $n = 0$ the claim is trivial. Now assume $n \neq 0$ and let $m = n\ div\ 2$ and $m' = n - 1 - m$ (and thus $m, m' < n$). From $bal\ n\ xs = (t, zs)$ we obtain l, r and ys such that $bal\ m\ xs = (l, ys)$, $bal\ m'\ (tl\ ys) = (r, zs)$ and $t = \langle l, hd\ ys, r\rangle$. Because $m < n \leq |xs|$ the induction hypothesis implies $xs = inorder\ l\ @\ ys \wedge |l| = m$ (1). This in turn implies $m' \leq |tl\ ys|$ and thus the induction hypothesis implies $tl\ ys = inorder\ r\ @\ zs \wedge |r| = m'$ (2). Properties (1) and (2) together with $t = \langle l, hd\ ys, r\rangle$ imply the claim $xs = inorder\ t\ @\ zs \wedge |t| = n$ because $ys \neq []$. □

The corresponding correctness properties of the derived functions are easy consequences:

$$\begin{aligned} n \leq |xs| \longrightarrow\ & inorder\ (bal_list\ n\ xs) = take\ n\ xs \\ & inorder\ (balance_list\ xs) = xs \\ n \leq |t| \longrightarrow\ & inorder\ (bal_tree\ n\ t) = take\ n\ (inorder\ t) \\ & inorder\ (balance_tree\ t) = inorder\ t \end{aligned}$$

To prove that *bal* returns an almost complete tree we determine its height and minimal height.

Lemma 4.10. $n \leq |xs| \wedge bal\ n\ xs = (t, zs) \longrightarrow h\ t = \lceil \lg\ (n + 1)\rceil$

Proof. The proof structure is the same as for Lemma 4.9 and we reuse the variable names introduced there. In the induction step we obtain the simplified induction hypothesese $h\ l = \lceil \lg\ (m + 1)\rceil$ and $h\ r = \lceil \lg\ (m' + 1)\rceil$. This leads to

$$\begin{aligned} & h\ t = max\ (h\ l)\ (h\ r) + 1 \\ & = h\ l + 1 && \text{because } m' \leq m \\ & = \lceil \lg\ (m + 1) + 1\rceil \\ & = \lceil \lg\ (n + 1)\rceil && \text{by (2.29)} \end{aligned}$$ □

The following complementary lemma is proved in the same way:

Lemma 4.11. $n \leq |xs| \wedge bal\ n\ xs = (t, zs) \longrightarrow mh\ t = \lfloor \lg\ (n + 1)\rfloor$

By definition of *acomplete* and because $\lceil x\rceil - \lfloor x\rfloor \leq 1$ we obtain that *bal* (and consequently the functions that build on it) returns an almost complete tree:

Corollary 4.12. $n \leq |xs| \wedge bal\ n\ xs = (t, ys) \longrightarrow acomplete\ t$

4.3.2 Exercises

Exercise 4.3. Find a formula B such that $acomplete\ \langle l,\ x,\ r\rangle = B$ where B may only contain the functions *acomplete*, *complete*, *height*, arithmetic and boolean operations, l and r, but in particular not *min_height* or *Node* $(= \langle _,\ _,\ _\rangle)$. Prove $acomplete\ \langle l,\ x,\ r\rangle = B$.

Exercise 4.4. Prove that the running time of function *bal* is linear in its first argument. (Isabelle hint: you need to refer to *bal* as *Balance.bal*.)

4.4 Augmented Trees

A tree of type $'a\ tree$ only stores elements of type $'a$. However, it is frequently necessary to store some additional information of type $'b$ in each node too, usually for efficiency reasons. Typical examples are:

- The size or the height of the tree. Because recomputing them requires traversing the whole tree.
- Lookup tables where each key of type $'a$ is associated with a value of type $'b$.

In this case we simply work with trees of type $('a \times 'b)\ tree$ and call them **augmented trees**. As a result we need to redefine a few functions that should ignore the additional information. For example, function *inorder*, when applied to an augmented tree, should return an $'a\ list$. Thus we redefine it in the obvious way:

```
inorder :: ('a × 'b) tree ⇒ 'a list
inorder ⟨⟩ = []
inorder ⟨l, (a, _), r⟩ = inorder l @ a # inorder r
```

Another example is $set_tree :: ('a \times 'b)\ tree \Rightarrow 'a\ set$. In general, if a function f is originally defined on type $'a\ tree$ but should ignore the $'b$-values in an $('a \times 'b)\ tree$ then we assume that there is a corresponding revised definition of f on augmented trees that focuses on the $'a$-values just like *inorder* above does. Of course functions that do not depend on the information in the nodes, e.g. size and height, stay unchanged.

Note that there are two alternative redefinitions of *inorder* (and similar functions): $map\ fst \circ Tree.inorder$ or $Tree.inorder \circ map_tree\ fst$ where *Tree.inorder* is the original function.

4.4.1 Maintaining Augmented Trees

Maintaining the $'b$-values in a tree can be hidden inside a suitable smart version of *Node* that has only a constant time overhead. Take the example of augmentation by size:

```
sz :: ('a × nat) tree ⇒ nat
sz ⟨⟩ = 0
sz ⟨_, (_, n), _⟩ = n

node_sz :: ('a × nat) tree ⇒ 'a ⇒ ('a × nat) tree ⇒ ('a × nat) tree
node_sz l a r = ⟨l, (a, sz l + sz r + 1), r⟩
```

A $('a \times nat)$ *tree* satisfies *invar_sz* if the size annotation of every node is computed from its children as specified in *node_sz*:

```
invar_sz :: ('a × nat) tree ⇒ bool
invar_sz ⟨⟩ = True
invar_sz ⟨l, (_, n), r⟩
= (n = sz l + sz r + 1 ∧ invar_sz l ∧ invar_sz r)
```

This property is preserved by *node_sz* (i.e. $invar_sz\ l \wedge invar_sz\ r \longrightarrow invar_sz\ (node_sz\ l\ a\ r)$) and it guarantees that *sz* returns the size:

$$invar_sz\ t \longrightarrow sz\ t = |t|$$

We can generalize this example easily. Assume we have a constant $zero :: 'b$ and a function $f :: 'b \Rightarrow 'a \Rightarrow 'b \Rightarrow 'b$ which we iterate over the tree:

```
F :: ('a × 'b) tree ⇒ 'b
F ⟨⟩ = zero
F ⟨l, (a, _), r⟩ = f (F l) a (F r)
```

This generalizes the definition of size. Let *node_f* compute the $'b$-value from the $'b$-value of its children via f:

```
b_val :: ('a × 'b) tree ⇒ 'b
b_val ⟨⟩ = zero
```

```
b_val ⟨_, (_, b), _⟩ = b

node_f :: ('a × 'b) tree ⇒ 'a ⇒ ('a × 'b) tree ⇒ ('a × 'b) tree
node_f l a r = ⟨l, (a, f (b_val l) a (b_val r)), r⟩
```

If all $'b$-values are computed as in *node_f*

```
invar_f :: ('a × 'b) tree ⇒ bool
invar_f ⟨⟩ = True
invar_f ⟨l, (a, b), r⟩
= (b = f (b_val l) a (b_val r) ∧ invar_f l ∧ invar_f r)
```

then all $'b$-values equal F: *invar_f t* $\longrightarrow$ *b_val t* $=$ *F t*.

4.4.2 Exercises

Exercise 4.5. Augment trees by a pair of a boolean and something else where the boolean indicates whether the tree is complete or not. Define *ch*, *node_ch* and *invar_ch* as in Section 4.4.1 and prove the following properties:

```
invar_ch t ⟶ ch t = (complete t, ?)
invar_ch l ∧ invar_ch r ⟶ invar_ch (node_ch l a r)
```

Exercise 4.6. Assume type $'a$ is of class *linorder* and augment each *Node* with the maximum value in that tree. Following Section 4.4.1 (but mind the *option* type!) define *mx* :: $('a \times 'b)$ *tree* $\Rightarrow$ $'b$ *option*, *node_mx* and *invar_mx* and prove

```
invar_mx t ⟶
mx t = (if t = ⟨⟩ then None else Some (Max (set_tree t)))
```

where *Max* is the predefined maximum operator on finite, non-empty sets.

5

Binary Search Trees

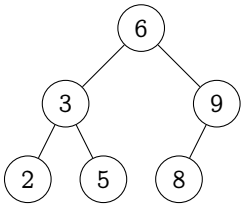
The purpose of this chapter is threefold: to introduce **binary search trees** (BSTs), to discuss their correctness proofs, and to provide a first example of an abstract data type, a notion discussed in more detail in the next chapter.

Search trees are a means for storing and accessing collections of elements efficiently. In particular they can support sets and maps. We concentrate on sets. We have already seen function *set_tree* that maps a tree to the set of its elements. This is an example of an **abstraction function** that maps concrete data structures to the abstract values that they represent.

BSTs require a linear ordering on the elements in the tree (as in Chapter Sorting). For each node, the elements in the left child are smaller than the root and the elements in the right child are bigger:

$$
\begin{aligned}
&bst :: ('a::linorder)\ tree \Rightarrow bool\\
&bst\ \langle\rangle = True\\
&bst\ \langle l,\ a,\ r\rangle\\
&= ((\forall x{\in}set_tree\ l.\ x < a) \wedge (\forall x{\in}set_tree\ r.\ a < x) \wedge bst\ l \wedge bst\ r)
\end{aligned}
$$

This is an example of a (coincidentally almost complete) BST:



It is obvious how to search for an element in a BST by comparing the element with the root and descending into one of the two children if you have

not found it yet. In the worst case this takes time proportional to the height of the tree. In later chapters we discuss a number of methods for ensuring that the height of the tree is logarithmic in its size. For now we ignore all efficiency considerations and permit our BSTs to degenerate.

Exercise 5.1. The above recursive definition of *bst* is not a direct translation of the description "For each node" given in the text. For a more direct translation define a function

$$nodes :: 'a\ tree \Rightarrow ('a\ tree \times 'a \times 'a\ tree)\ set$$

that collects all the nodes as triples $(l,\ a,\ r)$. Now define *bst_nodes* as $bst_nodes\ t = (\forall\,(l,a,r) \in nodes\ t.\ \ldots)$ and prove $bst_nodes\ t = bst\ t$.

5.1 Interfaces

Trees are concrete data types that provide the building blocks for realizing abstract data types like sets. The abstract type has a fixed interface, i.e. set of operations, through which the values of the abstract type can be manipulated. The interface hides all implementation detail. In the Search Tree part of the book we focus on the abstract type of sets with the following interface:

$$\begin{aligned}
&empty :: 's\\
&insert :: 'a \Rightarrow 's \Rightarrow 's\\
&delete :: 'a \Rightarrow 's \Rightarrow 's\\
&isin :: 's \Rightarrow 'a \Rightarrow bool
\end{aligned}$$

where *'s* is the type of sets of elements of type *'a*. Most of our implementations of sets will be based on variants of BSTs and will require a linear order on *'a*, but the general interface does not require this. The correctness of an implementation of this interface will be proved by relating it back to HOL's type *'a set* via an abstraction function, e.g. *set_tree*.

5.2 Implementing Sets via unbalanced BSTs

So far we have compared elements via $=$, $\leq$ and $<$. Now we switch to a comparator-based approach:

$$\textbf{datatype}\ cmp_val = LT \mid EQ \mid GT$$

$$\begin{aligned}
&cmp :: ('a :: linorder) \Rightarrow 'a \Rightarrow cmp_val\\
&cmp\ x\ y = (\textbf{if}\ x < y\ \textbf{then}\ LT\ \textbf{else if}\ x = y\ \textbf{then}\ EQ\ \textbf{else}\ GT)
\end{aligned}$$

We will frequently phrase algorithms in terms of *cmp*, LT, EQ and GT instead of $<$, $=$ and $>$. This leads to more symmetric code. If some type comes with its own primitive *cmp* function this can yield a speed-up over the above generic *cmp* function.

Below you find an implementation of the set interface in terms of BSTs. Functions *isin* and *insert* are self-explanatory. Deletion is more interesting.

```
empty :: 'a tree
empty = ⟨⟩

isin :: 'a tree ⇒ 'a ⇒ bool
isin ⟨⟩ _ = False
isin ⟨l, a, r⟩ x
= (case cmp x a of LT ⇒ isin l x | EQ ⇒ True | GT ⇒ isin r x)

insert :: 'a ⇒ 'a tree ⇒ 'a tree
insert x ⟨⟩ = ⟨⟨⟩, x, ⟨⟩⟩
insert x ⟨l, a, r⟩ = (case cmp x a of
                        LT ⇒ ⟨insert x l, a, r⟩ |
                        EQ ⇒ ⟨l, a, r⟩ |
                        GT ⇒ ⟨l, a, insert x r⟩)

delete :: 'a ⇒ 'a tree ⇒ 'a tree
delete x ⟨⟩ = ⟨⟩
delete x ⟨l, a, r⟩
= (case cmp x a of
    LT ⇒ ⟨delete x l, a, r⟩ |
    EQ ⇒ if r = ⟨⟩ then l
         else let (a', r') = split_min r in ⟨l, a', r'⟩ |
    GT ⇒ ⟨l, a, delete x r⟩)

split_min :: 'a tree ⇒ 'a × 'a tree
split_min ⟨l, a, r⟩
= (if l = ⟨⟩ then (a, r) else let (x, l') = split_min l in (x, ⟨l', a, r⟩))
```

5.2.1 Deletion

Function *delete* deletes a from $\langle l, a, r\rangle$ (where $r \neq \langle\rangle$) by replacing a with a' and r by r' where

a' is the leftmost (least) element of r, also called the inorder successor of a, r' is the remainder of r after removing a'.

We call this **deletion by replacing**. Of course one can also obtain a' as the inorder predecessor of a in l.

An alternative is to delete a from $\langle l, a, r\rangle$ by "joining" l and r:

```
delete2 :: 'a ⇒ 'a tree ⇒ 'a tree
delete2 _ ⟨⟩ = ⟨⟩
delete2 x ⟨l, a, r⟩ = (case cmp x a of
                        LT ⇒ ⟨delete2 x l, a, r⟩ |
                        EQ ⇒ join l r |
                        GT ⇒ ⟨l, a, delete2 x r⟩)

join :: 'a tree ⇒ 'a tree ⇒ 'a tree
join t ⟨⟩ = t
join ⟨⟩ t = t
join ⟨t1, a, t2⟩ ⟨t3, b, t4⟩
= (case join t2 t3 of
    ⟨⟩ ⇒ ⟨t1, a, ⟨⟨⟩, b, t4⟩⟩ |
    ⟨u2, x, u3⟩ ⇒ ⟨⟨t1, a, u2⟩, x, ⟨u3, b, t4⟩⟩)
```

We call this **deletion by joining**. The characteristic property of *join* is *inorder* (*join* l r) = *inorder* l @ *inorder* r.

The definition of *join* may appear needlessly complicated. Why not this much simpler version:

```
join0 t ⟨⟩ = t
join0 ⟨⟩ t = t
join0 ⟨t1, a, t2⟩ ⟨t3, b, t4⟩ = ⟨t1, a, ⟨join0 t2 t3, b, t4⟩⟩
```

Because with this version of *join*, deletion may almost double the height of the tree, in contrast to *join* and also deletion by replacing, where the height cannot increase:

Exercise 5.2. First prove that *join* behaves well:

$$h\ (join\ l\ r) \leq max\ (h\ l)\ (h\ r) + 1$$

Now show that *join*0 behaves badly: find an upper bound ub of h (*join*0 l r) such that ub is a function of $h\ l$ and $h\ r$. Prove h (*join*0 l r) $\leq ub$ and prove that ub is a tight upper bound if l and r are complete trees.

We focus on *delete*, deletion by replacing, in the rest of the chapter.

5.3 Correctness

Why is the above implementation correct? Roughly speaking, because the implementations of *empty*, *insert*, *delete* and *isin* on type *'a tree* simulate the behaviour of $\{\}$, $\cup$, $-$ and $\in$ on type *'a set*. Taking the abstraction function into account we can formulate the simulation precisely:

$$set_tree\ empty = \{\}$$
$$set_tree\ (insert\ x\ t) = set_tree\ t \cup \{x\}$$
$$set_tree\ (delete\ x\ t) = set_tree\ t - \{x\}$$
$$isin\ t\ x = (x \in set_tree\ t)$$

However, the implementation only works correctly on BSTs. Therefore we need to add the precondition *bst t* to all but the first property. But why are we permitted to assume this precondition? Only because *bst* is an **invariant** of this implementation: *bst* holds for *empty*, and both *insert* and *delete* preserve *bst*. Therefore every tree that can be manufactured through the interface is a BST. Of course this adds another set of proof obligations for correctness, **invariant preservation**:

$$bst\ empty$$
$$bst\ t \longrightarrow bst\ (insert\ x\ t)$$
$$bst\ t \longrightarrow bst\ (delete\ x\ t)$$

When looking at the abstract data type of sets from the user (or 'client') perspective, we would call the collection of all proof obligations for the correctness of an implementation the **specification** of the abstract type.

Exercise 5.3. Verify the implementation in Section 5.2 by showing all the proof obligations above, without the detour via sorted lists explained below.

5.4 Correctness Proofs

It turns out that direct proofs of the properties in the previous section are cumbersome — at least for *delete* and for proof assistants like Isabelle. Yet the correctness of the implementation is quite obvious to most (functional) programmers. Which is why most algorithm texts do not spend any time on functional correctness of search trees and concentrate on non-obvious structural properties that imply the logarithmic height of the trees — of course our simple BSTs do not guarantee the latter.

We will now present how the vague notion of "obvious" can be concretized and automated to such a degree that we do not need to discuss functional correctness of search tree implementations again in this book. This is because our

approach is quite generic: it works not only for the BSTs in this chapter but also for the more efficient variants discussed in later chapters. The remainder of this section can be skipped if one is not interested in proof automation.

5.4.1 The Idea

The key idea [54] is to express *bst* and *set_tree* via *inorder*:

$$bst\ t = sorted\ (inorder\ t) \quad \text{and} \quad set_tree\ t = set\ (inorder\ t)$$

where

```
sorted :: 'a list ⇒ bool
sorted [] = True
sorted [_] = True
sorted (x # y # zs) = (x < y ∧ sorted (y # zs))
```

Note that this is "sorted w.r.t. $(<)$" whereas in the chapter on sorting *sorted* was defined as "sorted w.r.t. $(\leq)$".

Instead of showing directly that BSTs implement sets, we show that they implement an intermediate specification based on lists (and later that the list-based specification implies the set-based one). We can assume that the lists are *sorted* because they are abstractions of BSTs. Insertion and deletion on sorted lists can be defined as follows:

```
ins_list :: 'a ⇒ 'a list ⇒ 'a list
ins_list x [] = [x]
ins_list x (a # xs)
= (if x < a then x # a # xs
   else if x = a then a # xs else a # ins_list x xs)

del_list :: 'a ⇒ 'a list ⇒ 'a list
del_list _ [] = []
del_list x (a # xs) = (if x = a then xs else a # del_list x xs)
```

The abstraction function from trees to lists is function *inorder*. The specification in Figure 5.1 expresses that *empty*, *insert*, *delete* and *isin* implement $[]$, *ins_list*, *del_list* and $\lambda xs\ x.\ x \in set\ xs$. One nice aspect of this specification is that it does not require us to prove invariant preservation explicitly: it follows from the fact (proved below) that *ins_list* and *del_list* preserve *sorted*.

$$inorder\ empty = []$$
$$sorted\ (inorder\ t) \longrightarrow inorder\ (insert\ x\ t) = ins_list\ x\ (inorder\ t)$$
$$sorted\ (inorder\ t) \longrightarrow inorder\ (delete\ x\ t) = del_list\ x\ (inorder\ t)$$
$$bst\ t \longrightarrow isin\ t\ x = (x \in set_tree\ t)$$

Fig. 5.1. List-based Specification of BSTs

5.4.2 BSTs Implement Sorted Lists — A Framework

We present a library of lemmas that automate the functional correctness proofs for the BSTs in this chapter and the more efficient variants in later chapters. This library is motivated by general considerations concerning the shape of formulas that arise during verification.

As a motivating example we examine how to prove

$$sorted\ (inorder\ t) \longrightarrow inorder\ (insert\ x\ t) = ins_list\ x\ (inorder\ t)$$

The proof is by induction on t and we consider the case $t = \langle l, a, r\rangle$ such that $x < a$. Ideally the proof looks like this:

$$\begin{aligned} &inorder\ (insert\ x\ t) = inorder\ (insert\ x\ l)\ @\ a\ \#\ inorder\ r \\ &= ins_list\ x\ (inorder\ l)\ @\ a\ \#\ inorder\ r \\ &= ins_list\ x\ (inorder\ l\ @\ a\ \#\ inorder\ r) = ins_list\ x\ t \end{aligned}$$

The first and last step are by definition, the second step by induction hypothesis, and the third step by lemmas in Figure 5.2: (5.1) rewrites the assumption $sorted\ (inorder\ t)$ to $sorted\ (inorder\ l\ @\ [a]) \wedge sorted\ (a\ \#\ inorder\ r)$, thus allowing (5.5) to rewrite the term $ins_list\ x\ (inorder\ l\ @\ a\ \#\ inorder\ r)$ to $ins_list\ x\ (inorder\ l)\ @\ a\ \#\ inorder\ r$.

The lemma library in Figure 5.2 helps to prove the properties in Figure 5.1. These proofs are by induction on t and lead to (possibly nested) tree constructor terms like $\langle\langle t_1, a_1, t_2\rangle, a_2, t_3\rangle$ where the t_i and a_i are variables. Evaluating *inorder* of such a tree leads to a list of the following form:

$$inorder\ t_1\ @\ a_1\ \#\ inorder\ t_2\ @\ a_2\ \#\ \ldots\ \#\ inorder\ t_n$$

Now we discuss the lemmas in Figure 5.2 that simplify the application of *sorted*, *ins_list* and *del_list* to such terms.

Terms of the form $sorted\ (xs_1\ @\ a_1\ \#\ xs_2\ @\ a_2\ \#\ \ldots\ \#\ xs_n)$ are decomposed into the following *basic* formulas

$$\begin{aligned} &sorted\ (xs\ @\ [a]) && \text{(simulating } \forall x{\in}set\ xs.\ x < a) \\ &sorted\ (a\ \#\ xs) && \text{(simulating } \forall x{\in}set\ xs.\ a < x) \\ &a < b \end{aligned}$$

by the rewrite rules (5.1)–(5.2). Lemmas (5.3)–(5.4) enable deductions from basic formulas.

$$sorted\ (xs\ @\ y\ \#\ ys) = (sorted\ (xs\ @\ [y]) \wedge sorted\ (y\ \#\ ys)) \tag{5.1}$$

$$\begin{aligned} &sorted\ (x\ \#\ xs\ @\ y\ \#\ ys) \\ &= (sorted\ (x\ \#\ xs) \wedge x < y \wedge sorted\ (xs\ @\ [y]) \wedge sorted\ (y\ \#\ ys)) \end{aligned} \tag{5.2}$$

$$sorted\ (x\ \#\ xs) \longrightarrow sorted\ xs \tag{5.3}$$

$$sorted\ (xs\ @\ [y]) \longrightarrow sorted\ xs \tag{5.4}$$

$$sorted\ (xs\ @\ [a]) \Longrightarrow ins_list\ x\ (xs\ @\ a\ \#\ ys) = \tag{5.5}$$
$$(\textbf{if}\ x < a\ \textbf{then}\ ins_list\ x\ xs\ @\ a\ \#\ ys\ \textbf{else}\ xs\ @\ ins_list\ x\ (a\ \#\ ys))$$

$$sorted\ (xs\ @\ a\ \#\ ys) \Longrightarrow del_list\ x\ (xs\ @\ a\ \#\ ys) = \tag{5.6}$$
$$(\textbf{if}\ x < a\ \textbf{then}\ del_list\ x\ xs\ @\ a\ \#\ ys\ \textbf{else}\ xs\ @\ del_list\ x\ (a\ \#\ ys))$$

$$sorted\ (x\ \#\ xs) = ((\forall\, y{\in}set\ xs.\ x < y) \wedge sorted\ xs) \tag{5.7}$$

$$sorted\ (xs\ @\ [x]) = (sorted\ xs \wedge (\forall\, y{\in}set\ xs.\ y < x)) \tag{5.8}$$

Fig. 5.2. Lemmas for *sorted*, *ins_list*, *del_list*

Terms of the form $ins_list\ x\ (xs_1\ @\ a_1\ \#\ xs_2\ @\ a_2\ \#\ \ldots\ \#\ xs_n)$ are rewritten with equation (5.5) (and the defining equations for *ins_list*) to push *ins_list* inwards. Terms of the form $del_list\ x\ (xs_1\ @\ a_1\ \#\ xs_2\ @\ a_2\ \#\ \ldots\ \#\ xs_n)$ are rewritten with equation (5.6) (and the defining equations for *del_list*) to push *del_list* inwards.

The *isin* property in Figure 5.1 can be proved with the help of (5.1), (5.7) and (5.8).

The lemmas in Figure 5.2 form the complete set of basic lemmas on which the automatic proofs of almost all search trees in the book rest; only splay trees (see Chapter 21) need additional lemmas.

5.4.3 Sorted Lists Implement Sets

It remains to be shown that the list-based specification (Figure 5.1) implies the set-based correctness properties in Section 5.3. Remember that $bst\ t = sorted\ (inorder\ t)$. The correctness properties

$$set_tree\ empty = \{\}$$

$$sorted\ (inorder\ t) \longrightarrow set_tree\ (insert\ x\ t) = set_tree\ t \cup \{x\}$$

$$sorted\ (inorder\ t) \longrightarrow set_tree\ (delete\ x\ t) = set_tree\ t - \{x\}$$

$$sorted\ (inorder\ t) \longrightarrow isin\ t\ x = (x \in set_tree\ t)$$

are a consequence of $set_tree\ t = set\ (inorder\ t)$ (a basic tree lemma), the properties in Figure 5.1 and the inductive correctness properties

$$set\ (ins_list\ x\ xs) = set\ xs \cup \{x\}$$

$$sorted\ xs \longrightarrow set\ (del_list\ x\ xs) = set\ xs - \{x\}$$

Preservation of the invariant

$$sorted\ (inorder\ empty)$$
$$sorted\ (inorder\ t) \longrightarrow sorted\ (inorder\ (insert\ x\ t))$$
$$sorted\ (inorder\ t) \longrightarrow sorted\ (inorder\ (delete\ x\ t))$$

are trivial or consequences of the properties in Figure 5.1 and the preservation of *sorted* by *ins_list* and *del_list*:

$$sorted\ xs \longrightarrow sorted\ (ins_list\ x\ xs)$$
$$sorted\ xs \longrightarrow sorted\ (del_list\ x\ xs)$$

This means in particular that preservation of $sorted \circ inorder$ is guaranteed for any implementation of *empty*, *insert* and *delete* that satisfies the list-based specification in Figure 5.1.

5.5 Interval Trees

In this section we study binary trees for representing a set of intervals, called **interval trees**. In addition to the usual insertion and deletion functions of standard BSTs, interval trees support a function for determining whether a given interval overlaps with some interval in the tree.

5.5.1 Augmented BSTs

The efficient implementation of the search for an overlapping interval relies on an additional piece of information in each node. Thus interval trees are another example of augmented trees as introduced in Section 4.4. We reuse the modified definitions of *set_tree* and *inorder* from that section. Moreover we use a slightly adjusted version of *isin*:

```
isin :: ('a × 'b) tree ⇒ 'a ⇒ bool
isin ⟨⟩ _ = False
isin ⟨l, (a, _), r⟩ x
= (case cmp x a of LT ⇒ isin l x | EQ ⇒ True | GT ⇒ isin r x)
```

This works for any kind of augmented BST, not just interval trees.

5.5.2 Intervals

An interval *'a ivl* is simply a pair of lower and upper bound, accessed by functions *low* and *high*, respectively. Intuitively, an interval represents the

closed set between *low* and *high*. The standard mathematical notation is $[l, h]$, the Isabelle notation is $\{l..h\}$. We restrict ourselves to non-empty intervals:

$$low\ p \leq high\ p$$

Type $'a$ can be any linearly ordered type with a minimum element $\bot$ (for example, the natural numbers or the real numbers extended with $-\infty$). Intervals can be linearly ordered by first comparing *low*, then comparing *high*. The definitions are as follows:

$$(x < y) = (low\ x < low\ y \vee low\ x = low\ y \wedge high\ x < high\ y)$$
$$(x \leq y) = (low\ x < low\ y \vee low\ x = low\ y \wedge high\ x \leq high\ y)$$

Two intervals overlap if they have at least one point in common:

$$overlap\ x\ y = (low\ y \leq high\ x \wedge low\ x \leq high\ y)$$

The readers should convince themselves that *overlap* does what it is supposed to do: $overlap\ x\ y = (\{low\ x..high\ x\} \cap \{low\ y..high\ y\} \neq \{\})$

We also define the concept of an interval overlapping with some interval in a set:

$$has_overlap\ S\ y = (\exists\, x{\in}S.\ overlap\ x\ y)$$

5.5.3 Interval Trees

An interval tree associates to each node a number *max_hi*, which records the maximum *high* value of all intervals in the subtrees. This value is updated during insert and delete operations, and will be crucial for enabling efficient determination of overlap with some interval in the tree.

type_synonym $'a\ ivl_tree = ('a\ ivl \times 'a)\ tree$

$max_hi :: 'a\ ivl_tree \Rightarrow 'a$
$max_hi\ \langle\rangle = \bot$
$max_hi\ \langle _, (_, m), _ \rangle = m$

If the *max_hi* value of every node in a tree agrees with *max*3

```
inv_max_hi :: 'a ivl_tree ⇒ bool
inv_max_hi ⟨⟩ = True
inv_max_hi ⟨l, (a, m), r⟩
= (m = max3 a (max_hi l) (max_hi r) ∧ inv_max_hi l ∧
   inv_max_hi r)

max3 :: 'a ivl ⇒ 'a ⇒ 'a ⇒ 'a
max3 a m n = max (high a) (max m n)
```

it follows by induction that *max_hi* is the maximum value of *high* in the tree and comes from some node in the tree:

Lemma 5.1. *inv_max_hi t ∧ a ∈ set_tree t ⟶ high a ≤ max_hi t*

Lemma 5.2. *inv_max_hi t ∧ t ≠ ⟨⟩ ⟶*
(∃ a∈set_tree t. high a = max_hi t)

5.5.4 Implementing Sets of Intervals via Interval Trees

Interval trees can implement sets of intervals via unbalanced BSTs as in Section 5.2. Of course *empty* = ⟨⟩. Function *isin* was already defined in Section 5.5.1 Insertion and deletion are also very close to the versions in Section 5.2, but the value of *max_hi* must be computed (by *max*3) for each new node. We follow Section 4.4 and introduce a smart constructor *node* for that purpose and replace ⟨*l*, *a*, *r*⟩ by *node l a r* (on the right-hand side):

```
node :: 'a ivl_tree ⇒ 'a ivl ⇒ 'a ivl_tree ⇒ 'a ivl_tree
node l a r = ⟨l, (a, max3 a (max_hi l) (max_hi r)), r⟩

insert :: 'a ivl ⇒ 'a ivl_tree ⇒ 'a ivl_tree
insert x ⟨⟩ = ⟨⟨⟩, (x, high x), ⟨⟩⟩
insert x ⟨l, (a, m), r⟩ = (case cmp x a of
                             LT ⇒ node (insert x l) a r |
                             EQ ⇒ ⟨l, (a, m), r⟩ |
                             GT ⇒ node l a (insert x r))

split_min :: 'a ivl_tree ⇒ 'a ivl × 'a ivl_tree
split_min ⟨l, (a, _), r⟩
= (if l = ⟨⟩ then (a, r)
```

```
    else let (x, l') = split_min l in (x, node l' a r))

delete :: 'a ivl ⇒ 'a ivl_tree ⇒ 'a ivl_tree
delete _ ⟨⟩ = ⟨⟩
delete x ⟨l, (a, _), r⟩
= (case cmp x a of
    LT ⇒ node (delete x l) a r |
    EQ ⇒ if r = ⟨⟩ then l else let (x, y) = split_min r in node l x y |
    GT ⇒ node l a (delete x r))
```

The correctness proofs for insertion and deletion cover two aspects. Functional correctness and preservation of the invariant $sorted \circ inorder$ (the BST property) are proved exactly as in Section 5.3 for ordinary BSTs. Preservation of the invariant *inv_max_hi* can be proved by a sequence of simple inductive properties. In the end the main correctness properties are

$$sorted\ (inorder\ t) \longrightarrow inorder\ (insert\ x\ t) = ins_list\ x\ (inorder\ t)$$
$$sorted\ (inorder\ t) \longrightarrow inorder\ (delete\ x\ t) = del_list\ x\ (inorder\ t)$$
$$inv_max_hi\ t \longrightarrow inv_max_hi\ (insert\ x\ t)$$
$$inv_max_hi\ t \longrightarrow inv_max_hi\ (delete\ x\ t)$$

Defining $invar\ t = (inv_max_hi\ t \wedge sorted\ (inorder\ t))$ we obtain the following top-level correctness corollaries:

$$invar\ s \longrightarrow set_tree\ (insert\ x\ s) = set_tree\ s \cup \{x\}$$
$$invar\ s \longrightarrow set_tree\ (delete\ x\ s) = set_tree\ s - \{x\}$$
$$invar\ s \longrightarrow invar\ (insert\ x\ s)$$
$$invar\ s \longrightarrow invar\ (delete\ x\ s)$$

The above insertion function allows overlapping intervals to be added into the tree and deletion supports only deletion of whole intervals. This is appropriate for the computational geometry application sketched below in Subsection 5.5.6. Other applications may require a different design.

5.5.5 Searching for an Overlapping Interval

The added functionality of interval trees over ordinary BSTs is function *search* that searches for an overlapping rather than identical interval:

```
search :: 'a ivl_tree ⇒ 'a ivl ⇒ bool
search ⟨⟩ _ = False
search ⟨l, (a, _), r⟩ x
= (if overlap x a then True
   else if l ≠ ⟨⟩ ∧ low x ≤ max_hi l then search l x else search r x)
```

The following theorem expresses the correctness of *search* assuming the same invariants as before; *bst t* would work just as well as *sorted* (*inorder t*).

Theorem 5.3. *inv_max_hi t* ∧ *sorted* (*inorder t*) ⟶ *search t x* = *has_overlap* (*set_tree t*) *x*

Proof. The result is clear when t is $\langle\rangle$. Now suppose t is in the form $\langle l, (a, m), r\rangle$, where m is the value of *max_hi* at root. If a overlaps with x, search returns *True* as expected. Otherwise, there are two cases.

- If $l \neq \langle\rangle$ and *low* $x \leq$ *max_hi* l, the search goes to the left child. If there is an interval in the left child overlapping with x, then the search returns *True* as expected. Otherwise, we show there is also no interval in the right child overlapping with x. Since $l \neq \langle\rangle$, Lemma 5.2 yields a node p in the left child such that *high* p = *max_hi* l. Since *low* $x \leq$ *max_hi* l, we have *low* $x \leq$ *high* p. Since p does not overlap with x, we must have *high* $x <$ *low* p. But then, for every interval *rp* in the right child, *low* $p \leq$ *low rp*, so that *high* $x <$ *low rp*, which implies that *rp* does not overlap with x.
- Now we consider the case where either $l = \langle\rangle$ or *max_hi* $l <$ *low* x. In this case, the search goes to the right. We show there is no interval in the left child that overlaps with x. This is clear if $l = \langle\rangle$. Otherwise, for each interval *lp*, we have *high lp* $\leq$ *max_hi* l by Lemma 5.1, so that *high lp* $<$ *low* x, which means *lp* does not overlap with x.

□

Exercise 5.4. Define a function that determines if a given point is in some interval in a given interval tree. Starting with

```
in_ivl :: 'a ⇒ 'a ivl ⇒ bool
in_ivl x iv = (low iv ≤ x ∧ x ≤ high iv)
```

write a recursive function

```
search1 :: 'a ivl_tree ⇒ 'a ⇒ bool
```

(without using *search*) such that *search*1 *x t* is *True* iff there is some interval *iv* in *t* such that *in_ivl x iv*. Prove

```
inv_max_hi t ∧ bst t ⟶ search1 t x = (∃ iv∈set_tree t. in_ivl x iv)
```

5.5.6 Application

While this section demonstrated how to augment an ordinary binary tree with intervals, any of the balanced binary trees (such as red-black tree) can be augmented in a similar manner. We leave this as exercises.

Interval trees have many applications in computational geometry. A basic example is as follows. Consider a set of rectangles whose sides are aligned to the x and y-axes. We wish to efficiently determine whether any pair of rectangles in the set intersect each other (i.e. sharing a point, including boundaries). This can be done using a "sweep line" algorithm as follows. For each rectangle $[x_l, x_h] \times [y_l, y_h]$, we create two events: insert interval $[x_l, x_h]$ at y-coordinate y_l and delete interval $[x_l, x_h]$ at y-coordinate y_h. Perform the events, starting from an empty interval tree, in ascending order of y-coordinates, with insertion events performed before deletion events. At each insertion, check whether the interval to be inserted overlaps with any of the existing intervals in the tree. If yes, we have found an intersection between two rectangles. If no overlap of intervals is detected throughout the process, then no pair of rectangles intersect. When using an interval tree based on a balanced binary tree, the time complexity of this procedure is $O(n \lg n)$, where n is the number of rectangles.

We refer to Cormen *et al.* [13, Section 14.3] for another exposition on interval trees and their applications. Interval trees, together with the application of finding rectangle intersection, have been formalized by Zhan [81].

6

Abstract Data Types

In the previous chapter we looked at a very specific example of an abstract data type, namely sets. In this chapter we consider abstract data types in general and in particular the model-oriented approach to the specification of abstract data types. This will lead to a generic format for such specifications.

6.1 Abstract Data Types

Abstract data types (ADTs) can be summarized by the following slogan:

$$\text{ADT} = \textit{interface} + \textit{specification}$$

where the interface lists the operations supported by the ADT and the specification describes the behaviour of these operations. For example, our set ADT has the following interface:

$$\begin{array}{l} \mathit{empty} :: 's \\ \mathit{insert} :: 'a \Rightarrow 's \Rightarrow 's \\ \mathit{delete} :: 'a \Rightarrow 's \Rightarrow 's \\ \mathit{isin} :: 's \Rightarrow 'a \Rightarrow \mathit{bool} \end{array}$$

The purpose of an ADT is to be able to write applications based on this ADT that will work with any implementation of the ADT. To this end one can prove properties of the application that are solely based on the specification of the ADT. That is, one can write generic algorithms and prove generic correctness theorems about them in the context of the ADT specification.

6.2 Model-Oriented Specifications

We follow the model-oriented style of specification [35]. In that style, an abstract type is specified by giving an abstract model for it. For simplicity we

assume that each ADT describes one **type of interest** T. In the set interface T is $'s$. This type T must be specified by some existing HOL type A, the abstract model. In the case of sets this is straightforward: the model for sets is simply the HOL type $'a\ set$. The motto is that T should behave like A. In order to bridge the gap between the two types, the specification needs an

- **abstraction function** $\alpha :: T \Rightarrow A$

that maps concrete values to their abstract counterparts. Moreover, in general only some elements of T represent elements of A. For example, in the set implementation in the previous chapter not all trees but only BSTs represent sets. Thus the specification should also take into account an

- **invariant** $invar :: T \Rightarrow bool$

Note that the abstraction function and the invariant are not part of the interface, but they are essential for specification and verification purposes.

As an example, the ADT of sets is shown in Figure 6.1 with suggestive keywords and a fixed mnemonic naming schema for the labels in the specification. This is the template for ADTs that we follow throughout the book.

ADT $Set =$

interface
$empty :: 's$
$insert :: 'a \Rightarrow 's \Rightarrow 's$
$delete :: 'a \Rightarrow 's \Rightarrow 's$
$isin :: 's \Rightarrow 'a \Rightarrow bool$

abstraction $set :: 's \Rightarrow 'a\ set$
invariant $invar :: 's \Rightarrow bool$

specification

$invar\ empty$	$(empty\text{-}inv)$
$set\ empty = \{\}$	$(empty)$
$invar\ s \longrightarrow invar\ (insert\ x\ s)$	$(insert\text{-}inv)$
$invar\ s \longrightarrow set(insert\ x\ s) = set\ s \cup \{x\}$	$(insert)$
$invar\ s \longrightarrow invar\ (delete\ x\ s)$	$(delete\text{-}inv)$
$invar\ s \longrightarrow set\ (delete\ x\ s) = set\ s - \{x\}$	$(delete)$
$invar\ s \longrightarrow isin\ s\ x = (x \in set\ s)$	$(isin)$

Fig. 6.1. ADT *Set*

We have intentionally refrained from showing the Isabelle formalization using so-called **locales** and have opted for a more intuitive textual format that is not Isabelle-specific, in accordance with the general philosophy of this book.

The actual Isabelle text can of course be found in the source files, and locales are explained in a dedicated manual [5].

We conclude this section by explaining what the specification of an arbitrary ADT looks like. We assume that for each function f of the interface there is a corresponding function f_A in the abstract model, i.e. defined on A. For a uniform treatment we extend α and *invar* to arbitrary types by setting $\alpha\ x = x$ and $invar\ x = True$ for all types other than T. Each function f of the interface gives rise to two properties in the specification: preservation of the invariant and simulation of f_A. The precondition is shared:

$$\begin{array}{ll} invar\ x_1 \wedge \ldots \wedge invar\ x_n \longrightarrow & \\ \quad invar(f\ x_1\ \ldots\ x_n) & (f\text{-}inv) \\ \quad \alpha(f\ x_1\ \ldots\ x_n) = f_A\ (\alpha\ x_1)\ \ldots\ (\alpha\ x_n) & (f) \end{array}$$

To understand how the specification of ADT *Set* is the result of this uniform schema one has to take two things into account:

- Precisely which abstract operations on type $'a\ set$ model the functions in the interface of the ADT *Set*? This correspondence is implicit in the specification: *empty* is modeled by $\{\}$, *insert* is modeled by $\lambda x\ s.\ s \cup \{x\}$, *delete* is modeled by $\lambda x\ s.\ s - \{x\}$ and *isin* is modeled by $\lambda s\ x.\ x \in s$.
- Because of the artificial extension of α and *invar* the above uniform format often collapses to something simpler where some α's and *invar*'s disappear.

6.3 Maps

An even more versatile data structure than sets are (efficient) maps from $'a$ to $'b$. In fact, sets can be viewed as maps from $'a$ to *bool*. Conversely, many data structures for sets also support maps, e.g. BSTs. Although, for simplicity, we focus on sets in this book, the ADT of maps should at least be introduced. It is shown in Figure 6.2. Type $'m$ is the type of maps from $'a$ to $'b$. The ADT *Map* is very similar to the ADT *Set* except that the abstraction function *lookup* is also part of the interface: it abstracts a map to a function of type $'a \Rightarrow 'b\ option$. This implies that the equations are between functions of that type. We use the function update notation (see Section 1.3) to explain *update* and *delete*: *update* is modeled by $\lambda m\ a\ b.\ m(a := b)$ and *delete* by $\lambda m\ a.\ m(a := None)$.

ADT *Map* =

interface
empty :: $'m$
update :: $'a \Rightarrow 'b \Rightarrow 'm \Rightarrow 'm$
delete :: $'a \Rightarrow 'm \Rightarrow 'm$
lookup :: $'m \Rightarrow 'a \Rightarrow 'b$ *option*

abstraction *lookup*
invariant *invar* :: $'m \Rightarrow$ *bool*

specification

invar empty	(*empty-inv*)
lookup empty $= (\lambda a.\ None)$	(*empty*)
invar $m \longrightarrow$ *invar* (*update a b m*)	(*update-inv*)
invar $m \longrightarrow$ *lookup* (*update a b m*) $=$ (*lookup m*)$(a := Some\ b)$	(*update*)
invar $m \longrightarrow$ *invar* (*delete a m*)	(*delete-inv*)
invar $m \longrightarrow$ *lookup* (*delete a m*) $=$ (*lookup m*)$(a := None)$	(*delete*)

Fig. 6.2. ADT *Map*

6.4 Implementing ADTs

An implementation of an ADT consists of definitions for all the functions in the interface. If you want to verify the correctness of the implementation, you also need to provide definitions for the abstraction function and the invariant. The latter two need not be executable unless they also occur in the interface and the implementation is meant to be executable. The abstraction function need not be surjective. For example, implementations of the ADT *Set* will normally only represent finite sets, e.g. by BSTs.

For Isabelle users: because ADTs are formalized as locales, an implementation of an ADT is an interpretation of the corresponding locale.

6.5 Exercises

Exercise 6.1. Modify the ADT *Set* specification by making *isin* the abstraction function (from $'s$ to $'a \Rightarrow$ *bool*). Follow the example of the ADT *Map* specification.

Exercise 6.2. In the ADT *Map* specification, following the general schema, there should be a property labeled (*lookup*), but it is missing. The reason is that given the correct abstract model of *lookup*, the equation becomes trivial: *invar* $m \longrightarrow$ *lookup* $m\ a =$ *lookup* $m\ a$. Why is that, which function models *lookup*?

Exercise 6.3. Implement ADT *Map* via unbalanced BSTs (like *Set* in Chapter 5) using augmented trees. Verify the implementation by proving all the correctness properties in the specification of ADT *Map* directly, without any detour via sorted lists as in Section 5.4.

7

2-3 Trees

This is the first in a series of chapters examining **balanced search trees** where the height of the tree is logarithmic in its size and which can therefore be searched in logarithmic time.

The most popular first example of balanced search trees are red-black trees. We start with **2-3 trees**, where nodes can have 2 or 3 children, because red-black trees are best understood as an implementation of (a variant of) 2-3 trees. We introduce red-black trees in the next chapter. The type of 2-3 trees is similar to binary trees but with an additional constructor *Node*3:

```
datatype 'a tree23 =
  Leaf |
  Node2 ('a tree23) 'a ('a tree23) |
  Node3 ('a tree23) 'a ('a tree23) 'a ('a tree23)
```

The familiar syntactic sugar is sprinkled on top:

$$\begin{aligned} \langle\rangle &\equiv \mathit{Leaf} \\ \langle l,\ a,\ r\rangle &\equiv \mathit{Node2}\ l\ a\ r \\ \langle l,\ a,\ m,\ b,\ r\rangle &\equiv \mathit{Node3}\ l\ a\ m\ b\ r \end{aligned}$$

The size, height and the completeness of a 2-3 tree are defined by adding an equation for *Node*3 to the corresponding definitions on binary trees:

$$|\langle l,\ _,\ m,\ _,\ r\rangle| = |l| + |m| + |r| + 1$$

```
h ⟨l, _, m, _, r⟩ = max (h l) (max (h m) (h r)) + 1

complete ⟨l, _, m, _, r⟩
= (h l = h m ∧ h m = h r ∧ complete l ∧ complete m ∧
   complete r)
```

A trivial induction yields $complete\ t \longrightarrow 2^{h\ t} \leq |t| + 1$: thus all operations on complete 2-3 trees have logarithmic complexity if they descend along a single branch and take constant time per node. This is the case and we will not discuss complexity in any more detail.

A nice property of 2-3 trees is that for every n there is a complete 2-3 tree of size n. As we will see below, completeness can be maintained under insertion and deletion in logarithmic time.

Exercise 7.1. Define a function $maxt :: nat \Rightarrow unit\ tree23$ that creates the tree with the largest number of nodes given the height of the tree. We use type *unit* because we are not interested in the elements in the tree. Prove $|maxt\ n| = (3^n - 1)\ div\ 2$ and that no tree of the given height can be larger: $|t| \leq (3^{h\ t} - 1)\ div\ 2$. Note that both subtraction and division on type *nat* can be tedious to work with. You may want to prove the two properties as corollaries of subtraction- and division-free properties. Alternatively, work with *real* instead of *nat* by replacing *div* by /.

7.1 Implementation of ADT *Set*

The implementation will maintain the usual ordering invariant and additionally completeness. When we speak of a 2-3 tree we will implicitly assume these two invariants now.

Searching a 2-3 tree is like searching a binary tree (see Section 5.2) but with one more defining equation:

```
isin ⟨l, a, m, b, r⟩ x
= (case cmp x a of LT ⇒ isin l x | EQ ⇒ True
   | GT ⇒ case cmp x b of LT ⇒ isin m x | EQ ⇒ True
           | GT ⇒ isin r x)
```

Insertion into a 2-3 tree must preserve the completeness invariant. Thus recursive calls must report back to the caller if the child has "overflown", i.e. increased in height. Therefore insertion returns a result of type $'a\ upI$:

```
datatype 'a upI = TI ('a tree23) | OF ('a tree23) 'a ('a tree23)
```

This is the idea: If insertion into t returns

$TI\ t'$ then t and t' should have the same height,
$OF\ l\ x\ r$ then t and l and r should have the same height.

The insertion functions are shown in Figure 7.1. The actual work is performed by the recursive function *ins*. The element to be inserted is propagated down to a leaf, which causes an overflow of the leaf. If an overflow is returned from a recursive call it can be absorbed into a *Node*2 but in a *Node*3 it causes another overflow. At the root of the tree, function *treeI* converts values of type *'a upI* back into trees:

```
treeI :: 'a upI ⇒ 'a tree23
treeI (TI t) = t
treeI (OF l a r) = ⟨l, a, r⟩
```

Deletion is dual. Recursive calls must report back to the caller if the child has "underflown", i.e. decreased in height. Therefore deletion returns a result of type *'a upD*:

```
datatype 'a upD = TD ('a tree23) | UF ('a tree23)
```

This is the idea: If deletion from t returns

$TD\ t'$ then t and t' should have the same height,
$UF\ t'$ then t should be one level higher than t'.

The main deletion functions are shown in Figure 7.2. The actual work is performed by the recursive function *del*. If the element to be deleted is in a child, the result of a recursive call is reintegrated into the node via the auxiliary functions *nodeij* from Figure 7.3 that create a node with i children where child j is given as an *'a upD* value. If the element to be deleted is in the node itself, a replacement is obtained and deleted from a child via *split_min*. At the root of the tree, values of type *'a upD* are converted back into trees:

```
treeD :: 'a upD ⇒ 'a tree23
treeD (TD t) = t
treeD (UF t) = t
```

```
insert x t = treeI (ins x t)

ins :: 'a ⇒ 'a tree23 ⇒ 'a upI
ins x ⟨⟩ = OF ⟨⟩ x ⟨⟩
ins x ⟨l, a, r⟩
= (case cmp x a of
    LT ⇒ case ins x l of TI l' ⇒ TI ⟨l', a, r⟩
           | OF l1 b l2 ⇒ TI ⟨l1, b, l2, a, r⟩
  | EQ ⇒ TI ⟨l, a, r⟩
  | GT ⇒ case ins x r of TI r' ⇒ TI ⟨l, a, r'⟩
           | OF r1 b r2 ⇒ TI ⟨l, a, r1, b, r2⟩)
ins x ⟨l, a, m, b, r⟩
= (case cmp x a of
    LT ⇒ case ins x l of TI l' ⇒ TI ⟨l', a, m, b, r⟩
           | OF l1 c l2 ⇒ OF ⟨l1, c, l2⟩ a ⟨m, b, r⟩
  | EQ ⇒ TI ⟨l, a, m, b, r⟩
  | GT ⇒ case cmp x b of
           LT ⇒ case ins x m of TI m' ⇒ TI ⟨l, a, m', b, r⟩
                  | OF m1 c m2 ⇒ OF ⟨l, a, m1⟩ c ⟨m2, b, r⟩
           | EQ ⇒ TI ⟨l, a, m, b, r⟩
           | GT ⇒ case ins x r of TI r' ⇒ TI ⟨l, a, m, b, r'⟩
                  | OF r1 c r2 ⇒ OF ⟨l, a, m⟩ b ⟨r1, c, r2⟩)
```

Fig. 7.1. Insertion into 2-3 tree

7.2 Preservation of Completeness

As explained in Section 5.4, we do not go into the automatic functional correctness proofs but concentrate on invariant preservation. To express the relationship between the height of a tree before and after insertion we define a height function *hI* on *'a upI* as follows:

```
hI (TI t) = h t
hI (OF l a r) = h l
```

The intuition is that *hI* is the height of the tree *before* insertion. A routine induction proves

$$complete\ t \longrightarrow complete\ (treeI\ (ins\ a\ t)) \wedge hI\ (ins\ a\ t) = h\ t$$

which implies by definition that

$$complete\ t \longrightarrow complete\ (insert\ a\ t)$$

```
delete x t = treeD (del x t)

del :: 'a ⇒ 'a tree23 ⇒ 'a upD
del x ⟨⟩ = TD ⟨⟩
del x ⟨⟨⟩, a, ⟨⟩⟩ = (if x = a then UF ⟨⟩ else TD ⟨⟨⟩, a, ⟨⟩⟩)
del x ⟨⟨⟩, a, ⟨⟩, b, ⟨⟩⟩
= TD (if x = a then ⟨⟨⟩, b, ⟨⟩⟩
        else if x = b then ⟨⟨⟩, a, ⟨⟩⟩ else ⟨⟨⟩, a, ⟨⟩, b, ⟨⟩⟩)
del x ⟨l, a, r⟩
= (case cmp x a of LT ⇒ node21 (del x l) a r
   | EQ ⇒ let (a', r') = split_min r in node22 l a' r'
   | GT ⇒ node22 l a (del x r))
del x ⟨l, a, m, b, r⟩
= (case cmp x a of LT ⇒ node31 (del x l) a m b r
   | EQ ⇒ let (a', m') = split_min m in node32 l a' m' b r
   | GT ⇒ case cmp x b of LT ⇒ node32 l a (del x m) b r
             | EQ ⇒ let (b', r') = split_min r in node33 l a m b' r'
             | GT ⇒ node33 l a m b (del x r))

split_min :: 'a tree23 ⇒ 'a × 'a upD
split_min ⟨⟨⟩, a, ⟨⟩⟩ = (a, UF ⟨⟩)
split_min ⟨⟨⟩, a, ⟨⟩, b, ⟨⟩⟩ = (a, TD ⟨⟨⟩, b, ⟨⟩⟩)
split_min ⟨l, a, r⟩ = (let (x, l') = split_min l in (x, node21 l' a r))
split_min ⟨l, a, m, b, r⟩
= (let (x, l') = split_min l in (x, node31 l' a m b r))
```

Fig. 7.2. Deletion from 2-3 tree: main functions

The fact that deletion preserves completeness can be proved by a sequence of small lemmas. To express the relationship between the height of a tree before and after deletion we define a height function *hD* on *'a upD* as follows:

```
hD (TD t) = h t
hD (UF t) = h t + 1
```

The intuition is that *hD* is the height of the tree *before* deletion. We now list a sequence of properties that build on each other and culminate in completeness preservation of *delete*:

```
complete r ∧ complete (treeD l') ∧ h r = hD l' ⟶
complete (treeD (node21 l' a r))
0 < h r ⟶ hD (node21 l' a r) = max (hD l') (h r) + 1
```

$$
\begin{aligned}
&node21 :: 'a\ upD \Rightarrow 'a \Rightarrow 'a\ tree23 \Rightarrow 'a\ upD\\
&node21\ (TD\ t_1)\ a\ t_2 = TD\ \langle t_1,\ a,\ t_2\rangle\\
&node21\ (UF\ t_1)\ a\ \langle t_2,\ b,\ t_3\rangle = UF\ \langle t_1,\ a,\ t_2,\ b,\ t_3\rangle\\
&node21\ (UF\ t_1)\ a\ \langle t_2,\ b,\ t_3,\ c,\ t_4\rangle = TD\ \langle\langle t_1,\ a,\ t_2\rangle,\ b,\ \langle t_3,\ c,\ t_4\rangle\rangle\\
\\
&node22 :: 'a\ tree23 \Rightarrow 'a \Rightarrow 'a\ upD \Rightarrow 'a\ upD\\
&node22\ t_1\ a\ (TD\ t_2) = TD\ \langle t_1,\ a,\ t_2\rangle\\
&node22\ \langle t_1,\ b,\ t_2\rangle\ a\ (UF\ t_3) = UF\ \langle t_1,\ b,\ t_2,\ a,\ t_3\rangle\\
&node22\ \langle t_1,\ b,\ t_2,\ c,\ t_3\rangle\ a\ (UF\ t_4) = TD\ \langle\langle t_1,\ b,\ t_2\rangle,\ c,\ \langle t_3,\ a,\ t_4\rangle\rangle\\
\\
&node31 :: 'a\ upD \Rightarrow 'a \Rightarrow 'a\ tree23 \Rightarrow 'a \Rightarrow 'a\ tree23 \Rightarrow 'a\ upD\\
&node31\ (TD\ t_1)\ a\ t_2\ b\ t_3 = TD\ \langle t_1,\ a,\ t_2,\ b,\ t_3\rangle\\
&node31\ (UF\ t_1)\ a\ \langle t_2,\ b,\ t_3\rangle\ c\ t_4 = TD\ \langle\langle t_1,\ a,\ t_2,\ b,\ t_3\rangle,\ c,\ t_4\rangle\\
&node31\ (UF\ t_1)\ a\ \langle t_2,\ b,\ t_3,\ c,\ t_4\rangle\ d\ t_5\\
&= TD\ \langle\langle t_1,\ a,\ t_2\rangle,\ b,\ \langle t_3,\ c,\ t_4\rangle,\ d,\ t_5\rangle\\
\\
&node32 :: 'a\ tree23 \Rightarrow 'a \Rightarrow 'a\ upD \Rightarrow 'a \Rightarrow 'a\ tree23 \Rightarrow 'a\ upD\\
&node32\ t_1\ a\ (TD\ t_2)\ b\ t_3 = TD\ \langle t_1,\ a,\ t_2,\ b,\ t_3\rangle\\
&node32\ t_1\ a\ (UF\ t_2)\ b\ \langle t_3,\ c,\ t_4\rangle = TD\ \langle t_1,\ a,\ \langle t_2,\ b,\ t_3,\ c,\ t_4\rangle\rangle\\
&node32\ t_1\ a\ (UF\ t_2)\ b\ \langle t_3,\ c,\ t_4,\ d,\ t_5\rangle\\
&= TD\ \langle t_1,\ a,\ \langle t_2,\ b,\ t_3\rangle,\ c,\ \langle t_4,\ d,\ t_5\rangle\rangle\\
\\
&node33 :: 'a\ tree23 \Rightarrow 'a \Rightarrow 'a\ tree23 \Rightarrow 'a \Rightarrow 'a\ upD \Rightarrow 'a\ upD\\
&node33\ l\ a\ m\ b\ (TD\ r) = TD\ \langle l,\ a,\ m,\ b,\ r\rangle\\
&node33\ t_1\ a\ \langle t_2,\ b,\ t_3\rangle\ c\ (UF\ t_4) = TD\ \langle t_1,\ a,\ \langle t_2,\ b,\ t_3,\ c,\ t_4\rangle\rangle\\
&node33\ t_1\ a\ \langle t_2,\ b,\ t_3,\ c,\ t_4\rangle\ d\ (UF\ t_5)\\
&= TD\ \langle t_1,\ a,\ \langle t_2,\ b,\ t_3\rangle,\ c,\ \langle t_4,\ d,\ t_5\rangle\rangle
\end{aligned}
$$

Fig. 7.3. Deletion from 2-3 tree: auxiliary functions

$$
\begin{aligned}
&split_min\ t = (x,\ t') \wedge 0 < h\ t \wedge complete\ t \longrightarrow hD\ t' = h\ t\\
&split_min\ t = (x,\ t') \wedge complete\ t \wedge 0 < h\ t \longrightarrow complete\ (treeD\ t')\\
&complete\ t \longrightarrow hD\ (del\ x\ t) = h\ t\\
&complete\ t \longrightarrow complete\ (treeD\ (del\ x\ t))\\
&complete\ t \longrightarrow complete\ (delete\ x\ t)
\end{aligned}
$$

For each property of *node*21 there are analogues properties for the other *node*ij functions which we omit.

7.3 Converting a List into a 2-3 Tree

The naive method of converting a list of elements into a 2-3 tree is to insert them one by one starting from the empty tree. However, that takes time $\Theta(n \lg n)$. This holds for any data structure where insertion takes time proportional to $\lg n$. In that case inserting n elements one by one takes time proportional to $\lg 1 + \cdots + \lg n = \lg(n!)$. Now $n! \le n^n$ implies $\lg(n!) \le n \lg n$. On the other hand, $n^n \le (n \cdot 1) \cdot ((n-1) \cdot 2) \cdots (1 \cdot n) = (n!)^2$ implies $\frac{1}{2} n \lg n \le \lg(n!)$. Thus $\lg(n!) \in \Theta(n \lg n)$ (which also follows from Stirling's formula). We have intentionally proved a Θ property because the O property is obvious but one might hope that $\lg 1 + \cdots + \lg n$ has a lower order of growth than $n \lg n$.

Luckily we can convert a whole list into a 2-3 tree in linear time. The bottom-up algorithm is particularly intuitive. It repeatedly passes over an alternating list $t_1, e_1, t_2, e_2, \ldots, t_k$ of trees and elements, combining trees and elements into new trees. Given elements $a_1, \ldots, a_n$ we start with the alternating list $\langle\rangle, a_1, \langle\rangle, a_2, \ldots, a_n, \langle\rangle$. On every pass over this list, we replace adjacent triples t, a, t' by $\langle t,\ a,\ t' \rangle$, possibly creating a 3-node instead of a 2-node at the end of the list. Once a single tree is left over, we terminate.

We define this type of alternating (and non-empty) list as a new data type:

```
datatype 'a tree23s = T ('a tree23) | TTs ('a tree23) 'a ('a tree23s)
```

The following examples demonstrate the encoding of alternating lists as terms of type *'a tree23s*:

Alternating list:	t_1	t_1, e_1, t_2	t_1, e_1, t_2, e_2, ts
Encoding:	$T\ t_1$	$TTs\ t_1\ e_1\ (T\ t_2)$	$TTs\ t_1\ e_1\ (TTs\ t_2\ e_2\ ts)$

We also need the following auxiliary functions:

```
len :: 'a tree23s ⇒ nat
len (T _) = 1
len (TTs _ _ ts) = len ts + 1

trees :: 'a tree23s ⇒ 'a tree23 set
trees (T t) = {t}
trees (TTs t _ ts) = {t} ∪ trees ts

inorder2 :: 'a tree23s ⇒ 'a list
inorder2 (T t) = inorder t
inorder2 (TTs t a ts) = inorder t @ a # inorder2 ts
```

Repeatedly passing over the alternating list until only a single tree remains is expressed by the following functions:

```
join_all :: 'a tree23s ⇒ 'a tree23
join_all (T t) = t
join_all ts = join_all (join_adj ts)

join_adj :: 'a tree23s ⇒ 'a tree23s
join_adj (TTs t1 a (T t2)) = T ⟨t1, a, t2⟩
join_adj (TTs t1 a (TTs t2 b (T t3))) = T ⟨t1, a, t2, b, t3⟩
join_adj (TTs t1 a (TTs t2 b ts)) = TTs ⟨t1, a, t2⟩ b (join_adj ts)
```

Note that *join_adj* is not and does not need to be defined on single trees. We express this precondition with an abbreviation:

$$not_T\ ts \equiv \forall t.\ ts \neq T\ t$$

Also note that *join_all* terminates only because *join_adj* shortens the list:

$$not_T\ ts \longrightarrow len\ (join_adj\ ts) < len\ ts$$

In fact, it reduces the length at least by a factor of 2:

$$not_T\ ts \longrightarrow len\ (join_adj\ ts) \leq len\ ts\ div\ 2 \tag{7.1}$$

The whole process starts with a list of alternating leaves and elements:

```
tree23_of_list :: 'a list ⇒ 'a tree23
tree23_of_list as = join_all (leaves as)

leaves :: 'a list ⇒ 'a tree23s
leaves [] = T ⟨⟩
leaves (a # as) = TTs ⟨⟩ a (leaves as)
```

7.3.1 Functional Correctness

Functional correctness is easily established. The *inorder* and the completeness properties are proved independently by the following inductive lemmas:

$$not_T\ ts \longrightarrow inorder2\ (join_adj\ ts) = inorder2\ ts$$

$$inorder\ (join_all\ ts) = inorder2\ ts$$

$$inorder\ (tree23_of_list\ as) = as$$

$$(\forall\, t{\in}\mathit{trees}\ ts.\ \mathit{complete}\ t \wedge h\ t = n) \wedge \mathit{not_T}\ ts \longrightarrow$$
$$(\forall\, t{\in}\mathit{trees}\ (\mathit{join_adj}\ ts).\ \mathit{complete}\ t \wedge h\ t = n + 1)$$

$$(\forall\, t{\in}\mathit{trees}\ ts.\ \mathit{complete}\ t \wedge h\ t = n) \longrightarrow \mathit{complete}\ (\mathit{join_all}\ ts)$$

$$t \in \mathit{trees}\ (\mathit{leaves}\ as) \longrightarrow \mathit{complete}\ t \wedge h\ t = 0$$

$$\mathit{complete}\ (\mathit{tree23_of_list}\ as)$$

7.3.2 Running Time Analysis

Why does the conversion take linear time? Because the first pass over an alternating list of length n takes n steps, the next pass $n/2$ steps, the next pass $n/4$ steps, etc, and this sums up to $2n$. The canonical time functions for the formal proof are shown in Appendix B.3. The following upper bound is easily proved by induction on the computation of $\mathit{join_adj}$:

$$\mathit{not_T}\ ts \longrightarrow T_{\mathit{join_adj}}\ ts \leq \mathit{len}\ ts\ \mathit{div}\ 2 \tag{7.2}$$

An upper bound $T_{\mathit{join_all}}\ ts \leq 2 \cdot \mathit{len}\ ts$ follows by induction on the computation of $\mathit{join_adj}$. We focus on the induction step:

$$\begin{aligned}
&T_{\mathit{join_all}}\ ts \\
&= T_{\mathit{join_adj}}\ ts + T_{\mathit{join_all}}\ (\mathit{join_adj}\ ts) + 1 \\
&\leq \mathit{len}\ ts\ \mathit{div}\ 2 + 2 \cdot \mathit{len}\ (\mathit{join_adj}\ ts) + 1 && \text{using (7.2) and IH} \\
&\leq \mathit{len}\ ts\ \mathit{div}\ 2 + 2 \cdot (\mathit{len}\ ts\ \mathit{div}\ 2) + 1 && \text{by (7.1)} \\
&\leq 2 \cdot \mathit{len}\ ts && \text{because } 1 \leq \mathit{len}\ ts
\end{aligned}$$

Now it is routine to derive

$$T_{\mathit{tree23_of_list}}\ as \leq 3 \cdot |as| + 4$$

Bibliographic Remarks

The invention of 2-3 trees is credited to Hopcroft in 1970 [13, p. 337]. Equational definitions were given by Hoffmann and O'Donnell [29] (only insertion) and Reade [66]. Our formalisation is based on the teaching material by Turbak [74] and the article by Hinze [28].

8

Red-Black Trees

Red-black trees are a popular implementation technique for BSTs: they guarantee logarithmic height just like 2-3 trees but the code is arguably simpler. The nodes are colored either red or black. Abstractly, red-black trees encode 2-3-4 trees where nodes have between 2 and 4 children. Each 2-3-4 node is encoded by a group of between 2 and 4 colored binary nodes as follows:

$$\begin{aligned}
\langle\rangle &\approx \langle\rangle \\
\langle A,a,B\rangle &\approx \langle A,a,B\rangle \\
\langle A,a,B,b,C\rangle &\approx \langle\langle A,a,B\rangle,b,C\rangle \text{ or } \langle A,a,\langle B,b,C\rangle\rangle \\
\langle A,a,B,b,C,c,D\rangle &\approx \langle\langle A,a,B\rangle,b,\langle C,c,D\rangle\rangle
\end{aligned}$$

Color expresses grouping: a black node is the root of a 2-3-4 node, a red node is part of a bigger 2-3-4 node. Thus a red-black tree needs to satisfy the following properties or invariants:

1. The root is black.
2. Every $\langle\rangle$ is considered black.
3. If a node is red, its children are black.
4. All paths from a node to a leaf have the same number of black nodes.

The final property expresses that the corresponding 2-3-4 tree is complete. The last two properties imply that the tree has logarithmic height (see below).

We implement red-black trees as binary trees augmented (see Section 4.4) with a color tag:

```
datatype color = Red | Black

type_synonym 'a rbt = ('a × color) tree
```

Some new syntactic sugar is sprinkled on top:

```
R l a r ≡ ⟨l, (a, Red), r⟩
B l a r ≡ ⟨l, (a, Black), r⟩
```

The following functions get and set the color of a node:

```
color :: 'a rbt ⇒ color
color ⟨⟩ = Black
color ⟨_, (_, c), _⟩ = c

paint :: color ⇒ 'a rbt ⇒ 'a rbt
paint _ ⟨⟩ = ⟨⟩
paint c ⟨l, (a, _), r⟩ = ⟨l, (a, c), r⟩
```

Note that the *color* of a leaf is by definition black.

8.1 Invariants

The above informal description of the red-black tree invariants is formalized as the predicate *rbt* which (for reasons of modularity) is split into a color and a height invariant *invc* and *invh*:

```
rbt :: 'a rbt ⇒ bool
rbt t = (invc t ∧ invh t ∧ color t = Black)
```

The color invariant expresses that red nodes must have black children:

```
invc :: 'a rbt ⇒ bool
invc ⟨⟩ = True
invc ⟨l, (_, c), r⟩
= ((c = Red ⟶ color l = Black ∧ color r = Black) ∧
   invc l ∧ invc r)
```

The height invariant expresses (via the **black height** *bh*) that all paths from the root to a leaf have the same number of black nodes:

```
invh :: 'a rbt ⇒ bool
invh ⟨⟩ = True
invh ⟨l, (_, _), r⟩ = (bh l = bh r ∧ invh l ∧ invh r)

bh :: 'a rbt ⇒ nat
bh ⟨⟩ = 0
bh ⟨l, (_, c), _⟩ = (if c = Black then bh l + 1 else bh l)
```

Note that although *bh* traverses only the left spine of the tree, the fact that *invh* traverses the complete tree ensures that all paths from the root to a leaf are considered. (See Exercise 8.2)

The split of the invariant into *invc* and *invh* improves modularity: frequently one can prove preservation of *invc* and *invh* separately, which facilitates proof search. For compactness we will mostly present the combined invariance properties.

8.1.1 Logarithmic Height

In a red-black tree, i.e. *rbt t*, every path from the root to a leaf has the same number of black nodes, and no such path has two red nodes in a row. Thus each leaf is at most twice as deep as any other leaf, and therefore $h\ t \le 2 \cdot \lg |t|_1$. The detailed proof starts with the key inductive relationship between height and black height

$$\begin{array}{l} invc\ t \wedge invh\ t \longrightarrow \\ h\ t \le 2 \cdot bh\ t + (\textbf{if}\ color\ t = Black\ \textbf{then}\ 0\ \textbf{else}\ 1) \end{array}$$

which has the easy corollary $rbt\ t \longrightarrow h\ t\ /\ 2 \le bh\ t$. Together with the easy inductive property

$$invc\ t \wedge invh\ t \longrightarrow 2^{bh\ t} \le |t|_1$$

this implies $2^{h\ t\ /\ 2} \le 2^{bh\ t} \le |t|_1$ and thus $h\ t \le 2 \cdot \lg |t|_1$ if *rbt t*.

8.2 Implementation of ADT *Set*

We implement sets by red-black trees that are also BSTs. As usual, we only discuss the proofs of preservation of the *rbt* invariant.

Function *isin* is implemented as for all augmented BSTs (see Section 5.5.1).

```
insert x t = paint Black (ins x t)

ins :: 'a ⇒ 'a rbt ⇒ 'a rbt
ins x ⟨⟩ = R ⟨⟩ x ⟨⟩
ins x (B l a r) = (case cmp x a of
                    LT ⇒ baliL (ins x l) a r |
                    EQ ⇒ B l a r |
                    GT ⇒ baliR l a (ins x r))
ins x (R l a r) = (case cmp x a of
                    LT ⇒ R (ins x l) a r |
                    EQ ⇒ R l a r |
                    GT ⇒ R l a (ins x r))

baliL :: 'a rbt ⇒ 'a ⇒ 'a rbt ⇒ 'a rbt
baliL (R (R t1 a t2) b t3) c t4 = R (B t1 a t2) b (B t3 c t4)
baliL (R t1 a (R t2 b t3)) c t4 = R (B t1 a t2) b (B t3 c t4)
baliL t1 a t2 = B t1 a t2

baliR :: 'a rbt ⇒ 'a ⇒ 'a rbt ⇒ 'a rbt
baliR t1 a (R t2 b (R t3 c t4)) = R (B t1 a t2) b (B t3 c t4)
baliR t1 a (R (R t2 b t3) c t4) = R (B t1 a t2) b (B t3 c t4)
baliR t1 a t2 = B t1 a t2
```

Fig. 8.1. Insertion into red-black tree

8.2.1 Insertion

Insertion is shown in Figure 8.1. The workhorse is function *ins*. It descends to the leaf where the element is inserted and it adjusts the colors on the way back up. The adjustment is performed by *baliL*/*baliR*. They combine arguments $l\ a\ r$ into a tree. If there is a red-red conflict in l/r, they rebalance and replace it by red-black. Inserting into a red node needs no immediate balancing because that will happen at the black node above it:

$$
\begin{aligned}
&ins\ 1\ (B\ (R\ \langle\rangle\ 0\ \langle\rangle)\ 2\ (R\ \langle\rangle\ 3\ \langle\rangle))\\
&= baliL\ (ins\ 1\ (R\ \langle\rangle\ 0\ \langle\rangle))\ 2\ (R\ \langle\rangle\ 3\ \langle\rangle)\\
&= baliL\ (R\ \langle\rangle\ 0\ (ins\ 1\ \langle\rangle))\ 2\ (R\ \langle\rangle\ 3\ \langle\rangle)\\
&= baliL\ (R\ \langle\rangle\ 0\ (R\ \langle\rangle\ 1\ \langle\rangle))\ 2\ (R\ \langle\rangle\ 3\ \langle\rangle)\\
&= R\ (B\ \langle\rangle\ 0\ \langle\rangle)\ 1\ (B\ \langle\rangle\ 2\ (R\ \langle\rangle\ 3\ \langle\rangle))
\end{aligned}
$$

Passing a red node up means an overflow occurred (as in 2-3 trees) that needs to be dealt with further up — at the latest at the very top where *insert* turns red into black.

Function *ins* preserves *invh* but not *invc*: it may return a tree with a red root node (see example above). However, once the root node is colored black, everything is fine again. Thus we introduce the weaker invariant *invc2* as an abbreviation:

$$invc2\ t \equiv invc\ (paint\ Black\ t)$$

It is easy to prove that *baliL* and *baliR* preserve *invh* and upgrade from *invc2* to *invc*:

$$invh\ l \wedge invh\ r \wedge invc2\ l \wedge invc\ r \wedge bh\ l = bh\ r \longrightarrow$$
$$invc\ (baliL\ l\ a\ r) \wedge invh\ (baliL\ l\ a\ r) \wedge bh\ (baliL\ l\ a\ r) = bh\ l + 1$$

$$invh\ l \wedge invh\ r \wedge invc\ l \wedge invc2\ r \wedge bh\ l = bh\ r \longrightarrow$$
$$invc\ (baliR\ l\ a\ r) \wedge invh\ (baliR\ l\ a\ r) \wedge bh\ (baliR\ l\ a\ r) = bh\ l + 1$$

Another easy induction yields

$$invc\ t \wedge invh\ t \longrightarrow$$
$$invc2\ (ins\ x\ t) \wedge (color\ t = Black \longrightarrow invc\ (ins\ x\ t)) \wedge$$
$$invh\ (ins\ x\ t) \wedge bh\ (ins\ x\ t) = bh\ t$$

The corollary $rbt\ t \longrightarrow rbt\ (insert\ x\ t)$ is immediate.

8.2.2 Deletion

Deletion from a red-black tree is shown in Figure 8.2. It follows the deletion-by-replacing approach (Section 5.2.1). The tricky bit is how to maintain the invariants. As before, intermediate trees may only satisfy the weaker invariant *invc2*. Functions *del* and *split_min* decrease the black height of a tree with a black root node and leave the black height unchanged otherwise. To see that this makes sense, consider deletion from a singleton black or red node. The case that the element to be removed is not in the black tree can be dealt with by coloring the root node red. These are the precise input/output relations:

Lemma 8.1. $split_min\ t = (x, t') \wedge t \neq \langle\rangle \wedge invh\ t \wedge invc\ t \longrightarrow$
$invh\ t' \wedge (color\ t = Red \longrightarrow bh\ t' = bh\ t \wedge invc\ t') \wedge$
$(color\ t = Black \longrightarrow bh\ t' = bh\ t - 1 \wedge invc2\ t')$

Lemma 8.2. $invh\ t \wedge invc\ t \wedge t' = del\ x\ t \longrightarrow$
$invh\ t' \wedge (color\ t = Red \longrightarrow bh\ t' = bh\ t \wedge invc\ t') \wedge$
$(color\ t = Black \longrightarrow bh\ t' = bh\ t - 1 \wedge invc2\ t')$

It is easy to see that the *del*-Lemma implies correctness of *delete*:

Corollary 8.3. $rbt\ t \longrightarrow rbt\ (delete\ x\ t)$

```
delete x t = paint Black (del x t)

del :: 'a ⇒ 'a rbt ⇒ 'a rbt
del _ ⟨⟩ = ⟨⟩
del x ⟨l, (a, _), r⟩
= (case cmp x a of
    LT ⇒ let l' = del x l
           in if l ≠ ⟨⟩ ∧ color l = Black then baldL l' a r else R l' a r |
    EQ ⇒ if r = ⟨⟩ then l
           else let (a', r') = split_min r
                  in if color r = Black then baldR l a' r' else R l a' r' |
    GT ⇒ let r' = del x r
           in if r ≠ ⟨⟩ ∧ color r = Black then baldR l a r' else R l a r')

split_min :: 'a rbt ⇒ 'a × 'a rbt
split_min ⟨l, (a, _), r⟩
= (if l = ⟨⟩ then (a, r)
    else let (x, l') = split_min l
         in (x, if color l = Black then baldL l' a r else R l' a r))

baldL :: 'a rbt ⇒ 'a ⇒ 'a rbt ⇒ 'a rbt
baldL (R t1 a t2) b t3 = R (B t1 a t2) b t3
baldL t1 a (B t2 b t3) = baliR t1 a (R t2 b t3)
baldL t1 a (R (B t2 b t3) c t4) = R (B t1 a t2) b (baliR t3 c (paint Red t4))
baldL t1 a t2 = R t1 a t2

baldR :: 'a rbt ⇒ 'a ⇒ 'a rbt ⇒ 'a rbt
baldR t1 a (R t2 b t3) = R t1 a (B t2 b t3)
baldR (B t1 a t2) b t3 = baliL (R t1 a t2) b t3
baldR (R t1 a (B t2 b t3)) c t4 = R (baliL (paint Red t1) a t2) b (B t3 c t4)
baldR t1 a t2 = R t1 a t2
```

Fig. 8.2. Deletion from red-black tree

The proofs of the two lemmas need the following precise characterizations of *baldL* and *baldR*, the counterparts of *baliL* and *baliR*:

Lemma 8.4. $invh\ l \wedge invh\ r \wedge bh\ l + 1 = bh\ r \wedge invc2\ l \wedge invc\ r \wedge$
$t' = baldL\ l\ a\ r \longrightarrow$
$invh\ t' \wedge bh\ t' = bh\ r \wedge invc2\ t' \wedge (color\ r = Black \longrightarrow invc\ t')$

Lemma 8.5. $invh\ l \wedge invh\ r \wedge bh\ l = bh\ r + 1 \wedge invc\ l \wedge invc2\ r \wedge$
$t' = baldR\ l\ a\ r \longrightarrow$
$invh\ t' \wedge bh\ t' = bh\ l \wedge invc2\ t' \wedge (color\ l = Black \longrightarrow invc\ t')$

The proofs of these lemmas are by case analyses over the defining equations using the characteristic properties of *baliL* and *baliR* given above.

Proof. Lemma 8.2 is proved by induction on the computation of $del\ x\ t$. The base case is trivial. In the induction step $t = \langle l, (a, c), r\rangle$. If $x < a$ then we distinguish three subcases. If $l = \langle\rangle$ the claim is trivial. Otherwise the claim follows from the IH: if $color\ l = Red$ then the claim follows directly, if $color\ l = Black$ then it follows with the help of Lemma 8.4 (with $l = del\ x\ l$). The case $a < x$ is dual and the case $x = a$ is similar (using Lemma 8.1). We do not show the details because they are tedious but routine. □

The proof of Lemma 8.1 is similar but simpler.

8.2.3 Deletion by Joining

As an alternative to deletion by replacement we also consider deletion by joining (see Section 5.2.1). The code for red-black trees is shown in Figure 8.3: compared to Figure 8.2, the EQ case of *del* has changed and *join* is new.

Invariant preservation is proved much like before except that instead of *split_min* we now have *join* to take care of. The characteristic lemma is proved by induction on the computation of *join*:

Lemma 8.6. $invh\ l \wedge invh\ r \wedge bh\ l = bh\ r \wedge invc\ l \wedge invc\ r \wedge$
$t' = join\ l\ r \longrightarrow$
$invh\ t' \wedge bh\ t' = bh\ l \wedge invc2\ t' \wedge$
$(color\ l = Black \wedge color\ r = Black \longrightarrow invc\ t')$

8.3 Exercises

Exercise 8.1. Show that the logarithmic height of red-black trees is already guaranteed by the color and height invariants:

$$invc\ t \wedge invh\ t \longrightarrow h\ t \leq 2 \cdot \lg |t|_1 + 2$$

Exercise 8.2. We already discussed informally why the definition of *invh* captures "all paths from the root to a leaf have the same number of black nodes" although *bh* only traverses the left spine. This exercises formalizes that discussion. The following function computes the set of black heights (number of black nodes) of all paths:

$bhs :: 'a\ rbt \Rightarrow nat\ set$
$bhs\ \langle\rangle = \{0\}$
$bhs\ \langle l, (_, c), r\rangle$
$= (\textbf{let}\ H = bhs\ l \cup bhs\ r\ \textbf{in}\ \textbf{if}\ c = Black\ \textbf{then}\ Suc\ `\ H\ \textbf{else}\ H)$

```
del :: 'a ⇒ 'a rbt ⇒ 'a rbt
del _ ⟨⟩ = ⟨⟩
del x ⟨l, (a, _), r⟩
= (case cmp x a of
    LT ⇒ if l ≠ ⟨⟩ ∧ color l = Black then baldL (del x l) a r
         else R (del x l) a r |
    EQ ⇒ join l r |
    GT ⇒ if r ≠ ⟨⟩ ∧ color r = Black then baldR l a (del x r)
         else R l a (del x r))

join :: 'a rbt ⇒ 'a rbt ⇒ 'a rbt
join ⟨⟩ t = t
join t ⟨⟩ = t
join (R t1 a t2) (R t3 c t4)
= (case join t2 t3 of
   R u2 b u3 ⇒ R (R t1 a u2) b (R u3 c t4) |
   t23 ⇒ R t1 a (R t23 c t4))
join (B t1 a t2) (B t3 c t4)
= (case join t2 t3 of
   R u2 b u3 ⇒ R (B t1 a u2) b (B u3 c t4) |
   t23 ⇒ baldL t1 a (R t23 c t4))
join t1 (R t2 a t3) = R (join t1 t2) a t3 |
join (R t1 a t2) t3 = R t1 a (join t2 t3)
```

Fig. 8.3. Deletion from red-black tree by combining children

where the infix operator (‘) is predefined as $f \mathbin{`} A = \{y \mid \exists x{\in}A.\ y = f\,x\}$. Prove $invh\ t \longleftrightarrow bhs\ t = \{bh\ t\}$. The $\longrightarrow$ direction should be easy, the other direction should need some lemmas.

Exercise 8.3. Following Section 7.3, define a linear time function *rbt_of_list* :: *'a list* ⇒ *'a rbt* and prove both *inorder* (*rbt_of_list as*) = *as* and *rbt* (*rbt_of_list as*).

Bibliographic Remarks

Red-Black trees were invented by Bayer [6] who called them “symmetric binary B-trees”. The red-black color convention was introduced by Guibas and Sedgewick [23] who studied their properties in greater depth. The first functional version of red-black trees (without deletion) is due to Okasaki [62] and everybody follows his code. A functional version of deletion was first given

by Kahrs [36][1] and Section 8.2.3 is based on it. Germane [22] presents a function for deletion by replacement that is quite different from the one in Section 8.2.2. Our starting point were Isabelle proofs by Reiter and Krauss (based on Kahrs). Other verifications of red-black trees are reported by Filliâtre and Letouzey [19] (using their own deletion function) and Appel [4] (based on Kahrs).

[1] The code for deletion is not in the article but can be retrieved from this URL: http://www.cs.ukc.ac.uk/people/staff/smk/redblack/rb.html

9

AVL Trees

The AVL tree [2] (named after its inventors Adelson-Velsky and Landis) is the granddaddy of efficient binary search trees. Its logarithmic height is maintained by rotating subtrees based on their height. For efficiency reasons the height of each subtree is stored in its root node. That is, the underlying data structure is a height-augmented tree (see Section 4.4):

```
type_synonym 'a tree_ht = ('a × nat) tree
```

Function *ht* extracts the height field and *node* is a smart constructor that sets the height field:

```
ht :: 'a tree_ht ⇒ nat
ht ⟨⟩ = 0
ht ⟨_, (_, n), _⟩ = n

node :: 'a tree_ht ⇒ 'a ⇒ 'a tree_ht ⇒ 'a tree_ht
node l a r = ⟨l, (a, max (ht l) (ht r) + 1), r⟩
```

An **AVL tree** is a tree that satisfies the AVL invariant: the height of the left and right child of any node differ by at most 1

```
avl :: 'a tree_ht ⇒ bool
avl ⟨⟩ = True
avl ⟨l, (_, n), r⟩
= (|int (h l) − int (h r)| ≤ 1 ∧
   n = max (h l) (h r) + 1 ∧ avl l ∧ avl r)
```

and the height field contains the correct value. The conversion function $int :: nat \Rightarrow int$ is required because on natural numbers $0 - n = 0$.

9.1 Logarithmic Height

AVL trees have logarithmic height. The key insight for the proof is that $M\ n$, the minimal number of leaves of an AVL tree of height n, satisfies the recurrence relation $M\ (n + 2) = M\ (n + 1) + M\ n$. Instead of formalizing this function M we prove directly that an AVL tree of height n has at least $fib\ (n + 2)$ leaves where fib is the Fibonacci function:

$$
\begin{aligned}
&fib :: nat \Rightarrow nat\\
&fib\ 0 = 0\\
&fib\ 1 = 1\\
&fib\ (n + 2) = fib\ (n + 1) + fib\ n
\end{aligned}
$$

Lemma 9.1. $avl\ t \longrightarrow fib\ (h\ t + 2) \leq |t|_1$

Proof. The proof is by induction on t. We focus on the induction step $t = \langle l, (a, n), r\rangle$ and assume $avl\ t$. Thus the IHs reduce to $fib\ (h\ l + 2) \leq |l|_1$ and $fib\ (h\ r + 2) \leq |r|_1$. We prove $fib\ (max\ (h\ l)\ (h\ r) + 3) \leq |l|_1 + |r|_1$, from which $avl\ t \longrightarrow fib\ (h\ t + 2) \leq |t|_1$ follows directly. There are two cases. We focus on $h\ l \geq h\ r$, $h\ l < h\ r$ is dual.

$$
\begin{aligned}
&fib\ (max\ (h\ l)\ (h\ r) + 3) = fib\ (h\ l + 3)\\
&= fib\ (h\ l + 2) + fib\ (h\ l + 1)\\
&\leq |l|_1 + fib\ (h\ l + 1) && \text{by } fib\ (h\ l + 2) \leq |l|_1\\
&\leq |l|_1 + |r|_1 && \text{by } fib\ (h\ r + 2) \leq |r|_1
\end{aligned}
$$

The last step is justified because $h\ l + 1 \leq h\ r + 2$ (which follows from $avl\ t$) and fib is monotone. □

Now we prove a well-known exponential lower bound for fib where $\varphi \equiv (1 + \sqrt{5})\ /\ 2$:

Lemma 9.2. $\varphi^n \leq fib\ (n + 2)$

Proof. The proof is by induction on n by fib computation induction. The case $n = 0$ is trivial and the case $n = 1$ is easy. Now consider the induction step:

$$\begin{aligned}
&fib\ (n + 2 + 2) = fib\ (n + 2 + 1) + fib\ (n + 2) \\
&\geq \varphi^{n+1} + \varphi^n && \text{by IHs} \\
&= (\varphi + 1) \cdot \varphi^n \\
&= \varphi^{n+2} && \text{because } \varphi + 1 = \varphi^2 \quad \square
\end{aligned}$$

Combining the two lemmas yields $avl\ t \longrightarrow \varphi^{h\ t} \leq |t|_1$ and thus

Corollary 9.3. $avl\ t \longrightarrow h\ t \leq 1\ /\ \lg\ \varphi \cdot \lg\ |t|_1$

That is, the height of an AVL tree is at most $1\ /\ \lg\ \varphi \approx 1.44$ times worse than the optimal $\lg\ |t|_1$.

9.2 Implementation of ADT *Set*

9.2.1 Insertion

Insertion follows the standard approach: insert the element as usual and reestablish the AVL invariant on the way back up.

```
insert :: 'a ⇒ 'a tree_ht ⇒ 'a tree_ht
insert x ⟨⟩ = ⟨⟨⟩, (x, 1), ⟨⟩⟩
insert x ⟨l, (a, n), r⟩ = (case cmp x a of
                             LT ⇒ balL (insert x l) a r |
                             EQ ⇒ ⟨l, (a, n), r⟩ |
                             GT ⇒ balR l a (insert x r))
```

Functions *balL*/*balR* readjust the tree after an insertion into the left/right child. The AVL invariant has been lost if the difference in height has become 2 — it cannot become more because the height can only increase by 1. Consider the definition of *balL* in Figure 9.1 (*balR* in Figure 9.2 is dual). If the AVL invariant has not been lost, i.e. if $ht\ AB \neq ht\ C + 2$, then we can just return the AVL tree $node\ AB\ c\ C$. But if $ht\ AB = ht\ C + 2$, we need to "rotate" the subtrees suitably. Clearly AB must be of the form $\langle A, (a, _), B\rangle$. There are two cases, which are illustrated in Figure 9.1. Rectangles denote trees. Rectangles of the same height denote trees of the same height. Rectangles with a +1 denote the additional level due to insertion of the new element.

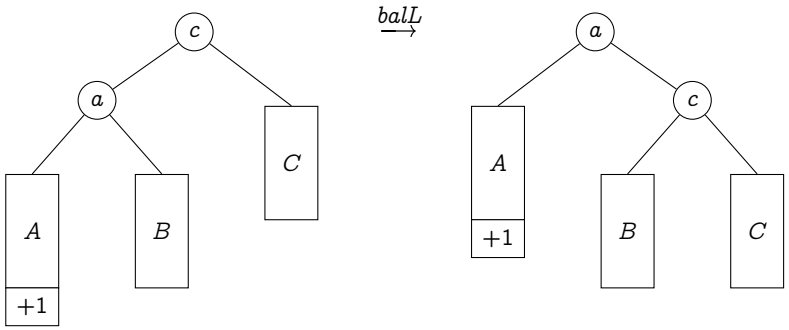
If $ht\ B \leq ht\ A$ then *balL* performs what is known as a single rotation.

If $ht\ A < ht\ B$ then B must be of the form $\langle B_1, (b, _), B_2\rangle$ (where either B_1 or B_2 has increased in height) and *balL* performs what is known as a double rotation.

It is easy to check that in both cases the tree on the right satisfies the AVL invariant.

```
balL :: 'a tree_ht ⇒ 'a ⇒ 'a tree_ht ⇒ 'a tree_ht
balL AB c C
= (if ht AB = ht C + 2
   then case AB of
        ⟨A, (a, x), B⟩ ⇒
          if ht B ≤ ht A then node A a (node B c C)
          else case B of
               ⟨B₁, (b, _), B₂⟩ ⇒ node (node A a B₁) b (node B₂ c C)
   else node AB c C)
```

Single rotation:



Double rotation:

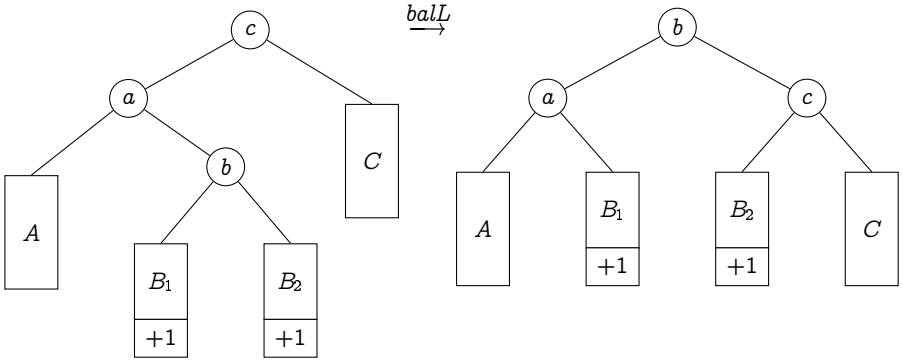


Fig. 9.1. Function *balL*

```
balR :: 'a tree_ht ⇒ 'a ⇒ 'a tree_ht ⇒ 'a tree_ht
balR A a BC
= (if ht BC = ht A + 2
   then case BC of
        ⟨B, (c, x), C⟩ ⇒
          if ht B ≤ ht C then node (node A a B) c C
          else case B of
               ⟨B₁, (b, _), B₂⟩ ⇒ node (node A a B₁) b (node B₂ c C)
   else node A a BC)
```

Fig. 9.2. Function *balR*

Preservation of *avl* by *insert* cannot be proved in isolation but needs to be proved simultaneously with how *insert* changes the height (because *avl* depends on the height and *insert* requires *avl* for correct behaviour):

Theorem 9.4. $avl\ t \longrightarrow avl\ (insert\ x\ t) \wedge h\ (insert\ x\ t) \in \{h\ t,\ h\ t + 1\}$

The proof is by induction on t followed by a complete case analysis (which Isabelle automates).

9.2.2 Deletion

```
delete :: 'a ⇒ 'a tree_ht ⇒ 'a tree_ht
delete _ ⟨⟩ = ⟨⟩
delete x ⟨l, (a, _), r⟩
= (case cmp x a of
   LT ⇒ balR (delete x l) a r |
   EQ ⇒ if l = ⟨⟩ then r else let (l', a') = split_max l in balR l' a' r |
   GT ⇒ balL l a (delete x r))

split_max :: 'a tree_ht ⇒ 'a tree_ht × 'a
split_max ⟨l, (a, _), r⟩
= (if r = ⟨⟩ then (l, a)
   else let (r', a') = split_max r in (balL l a r', a'))
```

Fig. 9.3. Deletion from AVL tree

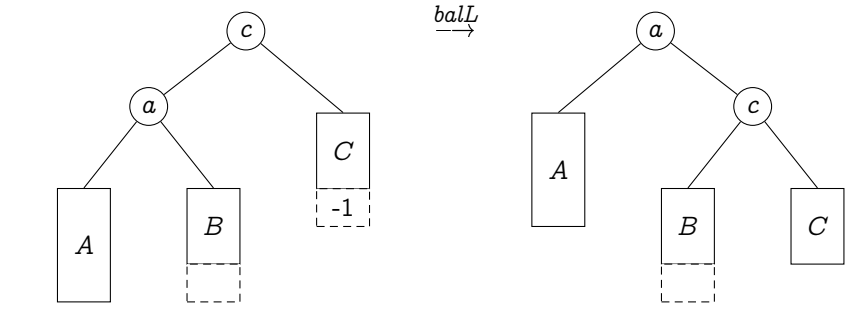
Figure 9.3 shows deletion-by-replacing (see 5.2.1). The recursive calls are dual to insertion: in terms of the difference in height, deletion of some element

from one child is the same as insertion of some element into the other child. Thus functions *balR*/*balL* can again be employed to restore the invariant.

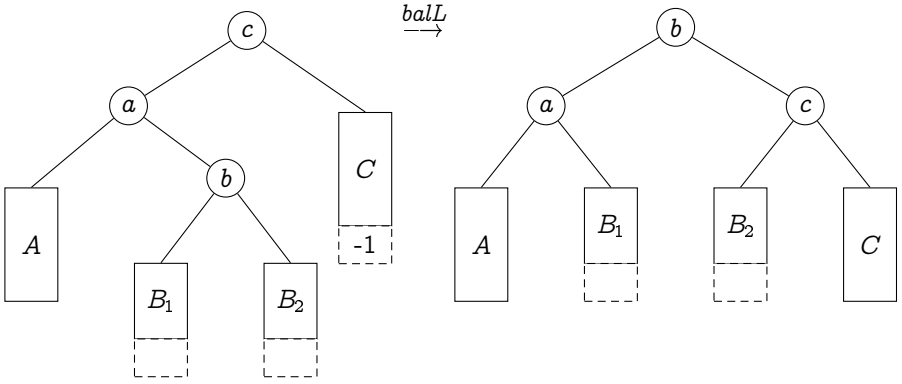
An element is deleted from a node by replacing it with the maximal element of the left child (the minimal element of the right child would work just as well). Function *split_max* performs that extraction and uses *balL* to restore the invariant after splitting an element off the right child.

The fact that *balR*/*balL* can be reused for deletion can be illustrated by drawing the corresponding rotation diagrams. We look at how the code for *balL* behaves when an element has been deleted from C. Dashed rectangles indicate a single additional level that may or may not be there. The label -1 indicates that the level has disappeared due to deletion.

Single rotation in *balL* after deletion in C:



Double rotation in *balL* after deletion in C:



At least one of B_1 and B_2 must have the same height as A.

Preservation of *avl* by *delete* can be proved in the same manner as for *insert* but we provide more of the details (partly because our Isabelle proof is less automatic). The following lemmas express that the auxiliary functions preserve *avl*:

$$avl\ l \wedge avl\ r \wedge h\ r - 1 \leq h\ l \wedge h\ l \leq h\ r + 2 \longrightarrow avl\ (balL\ l\ a\ r)$$

$$avl\ l \wedge avl\ r \wedge h\ l - 1 \leq h\ r \wedge h\ r \leq h\ l + 2 \longrightarrow avl\ (balR\ l\ a\ r)$$
$$avl\ t \wedge t \neq \langle\rangle \longrightarrow$$
$$avl\ (fst\ (split_max\ t)) \wedge$$
$$h\ t \in \{h\ (fst\ (split_max\ t)),\ h\ (fst\ (split_max\ t)) + 1\}$$

The first two are proved by the obvious cases analyses, the last one also requires induction.

As for *insert*, preservation of *avl* by *delete* needs to be proved simultaneously with how *delete* changes the height:

Theorem 9.5. $avl\ t \wedge t' = delete\ x\ t \longrightarrow avl\ t' \wedge h\ t \in \{h\ t',\ h\ t' + 1\}$

Proof. The proof is by induction on t followed by the case analyses dictated by the code for *delete*. We sketch the induction step. Let $t = \langle l, (a, n), r\rangle$ and $t' = delete\ x\ t$ and assume the IHs and $avl\ t$. The claim $avl\ t'$ follows from the preservation of *avl* by *balL*, *balR* and *split_max* as shown above. The claim $h\ t \in \{h\ t',\ h\ t' + 1\}$ follows directly from the definitions of *balL* and *balR*. □

9.3 Exercises

Exercise 9.1. The logarithmic height of AVL trees can be proved directly. Prove

$$avl\ t \wedge h\ t = n \longrightarrow 2^{n\ div\ 2} \leq |t|_1$$

by *fib* computation induction on n. This implies $avl\ t \longrightarrow h\ t \leq 2 \cdot \lg |t|_1$.

Exercise 9.2. Fibonacci trees are defined in analogy to Fibonacci numbers:

$$fibt :: nat \Rightarrow unit\ tree$$
$$fibt\ 0 = \langle\rangle$$
$$fibt\ 1 = \langle\langle\rangle, (), \langle\rangle\rangle$$
$$fibt\ (n + 2) = \langle fibt\ (n + 1), (), fibt\ n\rangle$$

We are only interested in the shape of these trees. Therefore the nodes just contain dummy unit values $()$. Hence we need to define the AVL invariant again for trees without annotations:

$$avl0 :: 'a\ tree \Rightarrow bool$$
$$avl0\ \langle\rangle = True$$
$$avl0\ \langle l, _, r\rangle = (|int\ (h\ l) - int\ (h\ r)| \leq 1 \wedge avl0\ l \wedge avl0\ r)$$

Prove the following properties of Fibonacci trees:

$$avl0\ (fibt\ n) \qquad |fibt\ n|_1 = fib\ (n + 2)$$

Conclude that the Fibonacci trees are minimal (w.r.t. their size) among all AVL trees of a given height:

$$avl\ t \longrightarrow |fibt\ (h\ t)|_1 \leq |t|_1$$

Exercise 9.3. Show that every almost complete tree is an AVL tree:

$$acomplete\ t \longrightarrow avl0\ t$$

As in the previous exercise we consider trees without height annotations.

Exercise 9.4. Generalize AVL trees to **height-balanced trees** where the condition

$$|int\ (h\ l) - int\ (h\ r)| \leq 1$$

in the invariant is replaced by

$$|int\ (h\ l) - int\ (h\ r)| \leq m$$

where $m \geq 1$ is some fixed integer. Modify the invariant and the insertion and deletion functions and prove that the latter fulfill the same correctness theorems as before. You do not need to prove the logarithmic height of height-balanced trees.

Exercise 9.5. Following Section 7.3, define a linear-time function $avl_of_list :: 'a\ list \Rightarrow 'a\ tree_ht$ and prove both $inorder\ (avl_of_list\ as) = as$ and $avl\ (avl_of_list\ as)$.

9.4 An Optimization

Instead of recording the height of the tree in each node, it suffices to record the **balance factor**, i.e. the difference in height of its two children. Rather than the three integers -1, 0 and 1 we utilize a new data type:

```
datatype bal = Lh | Bal | Rh

type_synonym 'a tree_bal = ('a × bal) tree
```

The names Lh and Rh stand for “left-heavy” and “right-heavy”. The AVL invariant for these trees reflect these names:

```
avl :: 'a tree_bal ⇒ bool
avl ⟨⟩ = True
avl ⟨l, (_, b), r⟩ = ((case b of
                        Lh ⇒ h l = h r + 1 |
                        Bal ⇒ h r = h l |
                        Rh ⇒ h r = h l + 1) ∧
                       avl l ∧ avl r)
```

The code for insertion (and deletion) is similar to the height-based version. The key difference is that the test if the AVL invariant as been lost cannot be based on the height anymore. We need to detect if the tree has increased in height upon insertion based on the balance factors. The key insight is that a height increase is coupled with a change from *Bal* to *Lh* or *Rh*. Except when we transition from ⟨⟩ to ⟨⟨⟩, (*a*, *Bal*), ⟨⟩⟩. This insight is encoded in the test *incr*:

```
is_bal :: 'a tree_bal ⇒ bool
is_bal ⟨_, (_, b), _⟩ = (b = Bal)

incr :: 'a tree_bal ⇒ 'b tree_bal ⇒ bool
incr t t' = (t = ⟨⟩ ∨ is_bal t ∧ ¬ is_bal t')
```

The test for a height increase compares the trees before and after insertion. Therefore it has been pulled out of the balance functions into insertion:

```
insert :: 'a ⇒ 'a tree_bal ⇒ 'a tree_bal
insert x ⟨⟩ = ⟨⟨⟩, (x, Bal), ⟨⟩⟩
insert x ⟨l, (a, b), r⟩
= (case cmp x a of
    LT ⇒ let l' = insert x l
         in if incr l l' then balL l' a b r else ⟨l', (a, b), r⟩ |
    EQ ⇒ ⟨l, (a, b), r⟩ |
    GT ⇒ let r' = insert x r
         in if incr r r' then balR l a b r' else ⟨l, (a, b), r'⟩)
```

The balance functions are shown in Figure 9.4. Function *rot2* implements double rotations. Function *balL* is called if the left child AB has increased in height. If the tree was *Lh* then single or double rotations are necessary

```
balL :: 'a tree_bal ⇒ 'a ⇒ bal ⇒ 'a tree_bal ⇒ 'a tree_bal
balL AB c bc C
= (case bc of
   Lh ⇒ case AB of
         ⟨A, (a, Lh), B⟩ ⇒ ⟨A, (a, Bal), ⟨B, (c, Bal), C⟩⟩ |
         ⟨A, (a, Bal), B⟩ ⇒ ⟨A, (a, Rh), ⟨B, (c, Lh), C⟩⟩ |
         ⟨A, (a, Rh), B⟩ ⇒ rot2 A a B c C |
   Bal ⇒ ⟨AB, (c, Lh), C⟩ |
   Rh ⇒ ⟨AB, (c, Bal), C⟩)

balR :: 'a tree_bal ⇒ 'a ⇒ bal ⇒ 'a tree_bal ⇒ 'a tree_bal
balR A a ba BC
= (case ba of
   Lh ⇒ ⟨A, (a, Bal), BC⟩ |
   Bal ⇒ ⟨A, (a, Rh), BC⟩ |
   Rh ⇒ case BC of
         ⟨B, (c, Lh), C⟩ ⇒ rot2 A a B c C |
         ⟨B, (c, Bal), C⟩ ⇒ ⟨⟨A, (a, Rh), B⟩, (c, Lh), C⟩ |
         ⟨B, (c, Rh), C⟩ ⇒ ⟨⟨A, (a, Bal), B⟩, (c, Bal), C⟩)

rot2 :: 'a tree_bal ⇒ 'a ⇒ 'a tree_bal ⇒ 'a ⇒ 'a tree_bal ⇒ 'a tree_bal
rot2 A a B c C
= (case B of
   ⟨B1, (b, bb), B2⟩ ⇒
     let b1 = if bb = Rh then Lh else Bal;
         b2 = if bb = Lh then Rh else Bal
     in ⟨⟨A, (a, b1), B1⟩, (b, Bal), ⟨B2, (c, b2), C⟩⟩)
```

Fig. 9.4. Functions *balL* and *balR*

to restore balance. Otherwise we simply need to adjust the balance factors. Function *balR* is dual to *balL*.

For deletion we need to test if the height has decreased and *decr* implements this test:

```
decr :: 'a tree_bal ⇒ 'b tree_bal ⇒ bool
decr t t' = (t ≠ ⟨⟩ ∧ (t' = ⟨⟩ ∨ ¬ is_bal t ∧ is_bal t'))
```

The functions *incr* and *decr* are almost dual except that *incr* implicitly assumes $t' \neq \langle\rangle$ because insertion is guaranteed to return a *Node*. Thus we could use *decr* instead of *incr* but not the other way around.

Deletion and *split_max* change in the same manner as insertion:

```
delete :: 'a ⇒ 'a tree_bal ⇒ 'a tree_bal

delete _ ⟨⟩ = ⟨⟩
delete x ⟨l, (a, ba), r⟩
= (case cmp x a of
    LT ⇒ let l' = delete x l
         in if decr l l' then balR l' a ba r else ⟨l', (a, ba), r⟩
  | EQ ⇒ if l = ⟨⟩ then r
         else let (l', a') = split_max l
              in if decr l l' then balR l' a' ba r
                 else ⟨l', (a', ba), r⟩
  | GT ⇒ let r' = delete x r
         in if decr r r' then balL l a ba r' else ⟨l, (a, ba), r'⟩)

split_max :: 'a tree_bal ⇒ 'a tree_bal × 'a

split_max ⟨l, (a, ba), r⟩
= (if r = ⟨⟩ then (l, a)
   else let (r', a') = split_max r;
            t' = if decr r r' then balL l a ba r' else ⟨l, (a, ba), r'⟩
        in (t', a'))
```

In the end we have the following correctness theorems:

Theorem 9.6. *avl t ∧ t' = insert x t ⟶*
avl t' ∧ h t' = h t + (if incr t t' then 1 else 0)

This theorem tells us not only that *avl* is preserved but also that *incr* indicates correctly if the height has increased or not.

Similarly for deletion and *decr*:

Theorem 9.7. *avl t ∧ t' = delete x t ⟶*
avl t' ∧ h t = h t' + (if decr t t' then 1 else 0)

The proofs of both theorems follow the standard pattern of induction followed by an exhaustive (automatic) cases analysis. The proof for *delete* requires an analogous lemma for *split_max*:

split_max t = (t', a) ∧ avl t ∧ t ≠ ⟨⟩ ⟶
avl t' ∧ h t = h t' + (if decr t t' then 1 else 0)

9.5 Exercises

Exercise 9.6. We map type *'a tree_bal* back to type *('a × nat) tree* called *'a tree_ht* in the beginning of the chapter:

$$debal :: 'a\ tree_bal \Rightarrow ('a \times nat)\ tree$$
$$debal\ \langle\rangle = \langle\rangle$$
$$debal\ \langle l, (a, _), r\rangle = \langle debal\ l, (a, max\ (h\ l)\ (h\ r) + 1), debal\ r\rangle$$

Prove that the AVL property is preserved: $avl\ t \longrightarrow avl_ht\ (debal\ t)$ where avl_ht is the avl predicate on type $'a\ tree_ht$ from the beginning of the chapter.

Define a function $debal2$ of the same type that traverses the tree only once and in particular does not use function h. Prove $avl\ t \longrightarrow debal2\ t = debal\ t$.

Exercise 9.7. To realize the full space savings potential of balance factors we encode them directly into the node constructors and work with the following special tree type:

datatype $'a\ tree4 = Leaf$
$\mid Lh\ ('a\ tree4)\ 'a\ ('a\ tree4)$
$\mid Bal\ ('a\ tree4)\ 'a\ ('a\ tree4)$
$\mid Rh\ ('a\ tree4)\ 'a\ ('a\ tree4)$

On this type define the AVL invariant, insertion, deletion and all necessary auxiliary functions. Prove theorems 9.6 and 9.7. Hint: modify the theory underlying Section 9.4.

10

Beyond Insert and Delete: $\cup$, $\cap$ and $-$

So far we looked almost exclusively at insertion and deletion of single elements, with the exception of the conversion of whole lists of elements into search trees (see Section 7.3 and Exercises 8.3 and 9.5). This chapter is dedicated to operations that combine two sets (implemented by search trees) by union, intersection and difference. We denote set difference by $-$ rather than $\setminus$.

Let us focus on set union for a moment and assume that insertion into a set of size s takes time proportional to $\lg s$. Consider two sets A and B of size m and n where $m \leq n$. The naive approach is to insert the elements from one set one by one into the other set. This takes time proportional to $\lg n + \cdots + \lg(n+m-1)$ or $\lg m + \cdots + \lg(m+n-1)$ depending on whether the smaller set is inserted into the larger one or the other way around. Of course the former sum is less than or equal to the latter sum. To estimate the growth of $\lg n + \cdots + \lg(n+m-1) = \lg(n \cdots (n+m-1))$ we can easily generalize the derivation of $\lg(n!) \in \Theta(n \lg n)$ in the initial paragraph of Section 7.3. The result is $\lg(n \cdots (n+m-1)) \in \Theta(m \lg n)$. That is, inserting m elements into an n element set one by one takes time $\Theta(m \lg n)$.

There is a second, possibly naive sounding algorithm for computing the union: flatten both trees to ordered lists (using function *inorder*2 from Exercise 4.1), merge both lists and convert the resulting list back into a suitably balanced search tree. All three steps take linear time. The last step is the only slightly nontrivial one but has been dealt with before (see Section 7.3 and Exercises 8.3 and 9.5). This algorithm takes time $O(m+n)$ which is significantly better than $O(m \lg n)$ if $m \approx n$ but significantly worse if $m \ll n$.

This chapter is concerned with a third approach that has the following salient features:

- Union, intersection and difference take time $O(m \lg(\frac{n}{m} + 1))$
- It works for a whole class of balanced trees, including AVL, red-black and weight-balanced trees.

- It is based on a single function for joining two balanced trees to form a new balanced tree.

We call it the **join approach**. It is easily and efficiently parallelizable, a property we will not explore here.

The join approach is at least as fast as the one-by-one approach: from $m+n \le mn$ it follows that $\frac{n}{m}+1 \le n$ (if $m, n \ge 2$). The join approach is also at least as fast as the tree-to-list-to-tree approach because $m+n = m(\frac{n}{m}+1)$ (if $m \ge 1$).

10.1 Specification of Union, Intersection and Difference

Before explaining the join approach we extend the ADT *Set* by three new functions *union*, *inter* and *diff*. The specification in Figure 10.1 is self-explanatory.

ADT *Set2* = *Set* +

interface

$union :: 's \Rightarrow 's \Rightarrow 's$
$inter :: 's \Rightarrow 's \Rightarrow 's$
$diff :: 's \Rightarrow 's \Rightarrow 's$

specification

$invar\ s_1 \wedge invar\ s_2 \longrightarrow invar\ (union\ s_1\ s_2)$	(*union-inv*)
$invar\ s_1 \wedge invar\ s_2 \longrightarrow set\ (union\ s_1\ s_2) = set\ s_1 \cup set\ s_2$	(*union*)
$invar\ s_1 \wedge invar\ s_2 \longrightarrow invar\ (inter\ s_1\ s_2)$	(*inter-inv*)
$invar\ s_1 \wedge invar\ s_2 \longrightarrow set\ (inter\ s_1\ s_2) = set\ s_1 \cap set\ s_2$	(*inter*)
$invar\ s_1 \wedge invar\ s_2 \longrightarrow invar\ (diff\ s_1\ s_2)$	(*diff-inv*)
$invar\ s_1 \wedge invar\ s_2 \longrightarrow set\ (diff\ s_1\ s_2) = set\ s_1 - set\ s_2$	(*diff*)

Fig. 10.1. ADT *Set2*

10.2 Just Join

Now we come to the heart of the matter, the definition of union, intersection and difference in terms of a single function *join*. We promised that the algorithms would be generic across a range of balanced trees. Thus we assume that we operate on augmented trees of type $('a \times 'b)\ tree$ where $'a$ is the type of the elements and $'b$ is the balancing information (which we can ignore

here). This enables us to formulate the algorithms via pattern-matching. A more generic approach is the subject of Exercise 10.1.

The whole section is parameterized by the join function and an invariant:

$$join :: ('a \times 'b)\ tree \Rightarrow 'a \Rightarrow ('a \times 'b)\ tree \Rightarrow ('a \times 'b)\ tree$$
$$inv :: ('a \times 'b)\ tree \Rightarrow bool$$

Function *inv* is meant to take care of the balancedness property only, not the BST property. Functions *join* and *inv* are specified with the help of the standard tree functions *set_tree* and *bst* in Figure 10.2. With respect to the

$$set_tree\ (join\ l\ a\ r) = set_tree\ l \cup \{a\} \cup set_tree\ r \tag{10.1}$$
$$bst\ \langle l, (a, _), r\rangle \longrightarrow bst\ (join\ l\ a\ r) \tag{10.2}$$
$$inv\ \langle\rangle$$
$$inv\ l \wedge inv\ r \longrightarrow inv\ (join\ l\ a\ r) \tag{10.3}$$
$$inv\ \langle l, (_, _), r\rangle \longrightarrow inv\ l \wedge inv\ r \tag{10.4}$$

Fig. 10.2. Specification of *join* and *inv*

set of elements, *join* must behave like union. But it need only return a BST if both trees are BSTs and the element a lies in between the elements of the two trees, i.e. if $bst\ \langle l, (a, _), r\rangle$. The structural invariant *inv* must be preserved by formation and destruction of trees. Thus we can see *join* as a smart constructor that builds a balanced tree.

To define union and friends we need a number of simple auxiliary functions shown in Figure 10.3. Function *split_min* decomposes a tree into its leftmost (minimal) element and the remaining tree; the remaining tree is reassembled via *join*, thus preserving *inv*. Function *join2* is reduced to *join* with the help of *split_min*. Function *split* splits a BST w.r.t. a given element a into a triple (l, b, r) such that l contains the elements less than a, r contains the elements greater than a, and b is true iff a was in the input tree.

Although insertion and deletion could be defined by means of union and difference, we can define them directly from the auxiliary functions:

$$insert :: 'a \Rightarrow ('a \times 'b)\ tree \Rightarrow ('a \times 'b)\ tree$$
$$insert\ x\ t = (\textbf{let}\ (l, b, r) = split\ t\ x\ \textbf{in}\ join\ l\ x\ r)$$

$$delete :: 'a \Rightarrow ('a \times 'b)\ tree \Rightarrow ('a \times 'b)\ tree$$
$$delete\ x\ t = (\textbf{let}\ (l, b, r) = split\ t\ x\ \textbf{in}\ join2\ l\ r)$$

The efficiency can be improved a little by taking the returned b into account.

```
split_min :: ('a × 'b) tree ⇒ 'a × ('a × 'b) tree
split_min ⟨l, (a, _), r⟩
= (if l = ⟨⟩ then (a, r)
   else let (m, l') = split_min l in (m, join l' a r))

join2 :: ('a × 'b) tree ⇒ ('a × 'b) tree ⇒ ('a × 'b) tree
join2 l r
= (if r = ⟨⟩ then l else let (m, r') = split_min r in join l m r')

split :: ('a × 'b) tree ⇒ 'a ⇒ ('a × 'b) tree × bool × ('a × 'b) tree
split ⟨⟩ _ = (⟨⟩, False, ⟨⟩)
split ⟨l, (a, _), r⟩ x
= (case cmp x a of
   LT ⇒ let (l1, b, l2) = split l x in (l1, b, join l2 a r) |
   EQ ⇒ (l, True, r) |
   GT ⇒ let (r1, b, r2) = split r x in (join l a r1, b, r2))
```

Fig. 10.3. Auxiliary functions

But we have bigger functions to fry: union, intersection and difference. They are shown in Figure 10.4. All three are divide-and-conquer algorithms that follow the same schema: both input trees are split at an element a (by construction or explicitly), the algorithm is applied recursively to the two trees of the elements below a and to the two trees of the elements above a, and the two results are suitably joined.

10.2.1 Correctness

We need to prove that *union*, *inter* and *diff* satisfy the specification in Figure 10.1 where $set = set_tree$ and $invar\ t = inv\ t \wedge bst\ t$. That is, for each function we show its set-theoretic property and preservation of *inv* and *bst* using the assumptions in Figure 10.2. Most of the proofs in this section are obvious and automatic inductions and we do not discuss them.

First we need to prove suitable properties of the auxiliary functions *split_min*, *join2* and *split*:

$$split_min\ t = (m, t') \wedge t \neq \langle\rangle \longrightarrow m \in set_tree\ t \wedge set_tree\ t = \{m\} \cup set_tree\ t'$$

$$split_min\ t = (m, t') \wedge bst\ t \wedge t \neq \langle\rangle \longrightarrow bst\ t' \wedge (\forall x \in set_tree\ t'.\ m < x)$$

$$split_min\ t = (m, t') \wedge inv\ t \wedge t \neq \langle\rangle \longrightarrow inv\ t'$$

```
union :: ('a × 'b) tree ⇒ ('a × 'b) tree ⇒ ('a × 'b) tree
union ⟨⟩ t = t
union t ⟨⟩ = t
union ⟨l1, (a, _), r1⟩ t2
= (let (l2, b2, r2) = split t2 a
   in join (union l1 l2) a (union r1 r2))

inter :: ('a × 'b) tree ⇒ ('a × 'b) tree ⇒ ('a × 'b) tree
inter ⟨⟩ t = ⟨⟩
inter t ⟨⟩ = ⟨⟩
inter ⟨l1, (a, _), r1⟩ t2
= (let (l2, b2, r2) = split t2 a;
       l' = inter l1 l2; r' = inter r1 r2
   in if b2 then join l' a r' else join2 l' r')

diff :: ('a × 'b) tree ⇒ ('a × 'b) tree ⇒ ('a × 'b) tree
diff ⟨⟩ t = ⟨⟩
diff t ⟨⟩ = t
diff t1 ⟨l2, (a, _), r2⟩
= (let (l1, b1, r1) = split t1 a
   in join2 (diff l1 l2) (diff r1 r2))
```

Fig. 10.4. Union, intersection and difference

$$set_tree\ (join2\ l\ r) = set_tree\ l \cup set_tree\ r \tag{10.5}$$

$$\begin{aligned}
&bst\ l \wedge bst\ r \wedge (\forall x{\in}set_tree\ l.\ \forall y{\in}set_tree\ r.\ x < y) \longrightarrow \\
&bst\ (join2\ l\ r)
\end{aligned}$$

$$inv\ l \wedge inv\ r \longrightarrow inv\ (join2\ l\ r)$$

$$\begin{aligned}
&split\ t\ x = (l,\ b,\ r) \wedge bst\ t \longrightarrow \\
&set_tree\ l = \{a \in set_tree\ t \mid a < x\} \wedge \\
&set_tree\ r = \{a \in set_tree\ t \mid x < a\} \wedge \\
&b = (x \in set_tree\ t) \wedge bst\ l \wedge bst\ r
\end{aligned} \tag{10.6}$$

$$split\ t\ x = (l,\ b,\ r) \wedge inv\ t \longrightarrow inv\ l \wedge inv\ r$$

The correctness properties of *insert* and *delete* are trivial consequences and are not shown. We move on to *union*. Its correctness properties are concretizations of the properties (*union*) and (*union-inv*) in Figure 10.1:

$$bst\ t_2 \longrightarrow set_tree\ (union\ t_1\ t_2) = set_tree\ t_1 \cup set_tree\ t_2$$

$$bst\ t_1 \wedge bst\ t_2 \longrightarrow bst\ (union\ t_1\ t_2)$$

$$inv\ t_1 \wedge inv\ t_2 \longrightarrow inv\ (union\ t_1\ t_2)$$

All three *union* properties are proved by computation induction. The first property follows easily from assumption (10.1) and (10.6). The assumption $bst\ t_2$ (but not $bst\ t_1$) is required because t_2 is split and (10.6) requires *bst*. Preservation of *bst* follows from assumption (10.2) with the help of the first *union* property and the preservation of *bst* by *split*. Preservation of *inv* follows from assumptions (10.3) and (10.4) with the help of the preservation of *inv* by *split*.

The correctness properties of *inter* look similar:

$$bst\ t_1 \wedge bst\ t_2 \longrightarrow set_tree\ (inter\ t_1\ t_2) = set_tree\ t_1 \cap set_tree\ t_2$$
$$bst\ t_1 \wedge bst\ t_2 \longrightarrow bst\ (inter\ t_1\ t_2)$$
$$inv\ t_1 \wedge inv\ t_2 \longrightarrow inv\ (inter\ t_1\ t_2)$$

The proof of the preservation properties are automatic but the proof of the *set_tree* property is more involved than the corresponding proof for *union* and we take a closer look at the induction. We focus on the case $t_1 = \langle l_1, (a, _), r_1\rangle$ and $t_2 \neq \langle\rangle$. Let $L_1 = set_tree\ l_1$ and $R_1 = set_tree\ r_1$. Let $(l_2, b, r_2) = split\ t_2\ a$, $L_2 = set_tree\ l_2$, $R_2 = set_tree\ r_2$ and $A =$ (**if** b **then** $\{a\}$ **else** $\{\}$). The separation properties

$$a \notin L_1 \cup R_1 \quad a \notin L_2 \cup R_2$$
$$L_2 \cap R_2 = \{\} \quad L_1 \cap R_2 = \{\} \quad L_2 \cap R_1 = \{\}$$

follow from $bst\ t_1$, $bst\ t_2$ and (10.6). Now for the main proof:

$$\begin{aligned}
&set_tree\ t_1 \cap set_tree\ t_2 \\
&= (L_1 \cup R_1 \cup \{a\}) \cap (L_2 \cup R_2 \cup A) && \text{by (10.6), } bst\ t_2 \\
&= L_1 \cap L_2 \cup R_1 \cap R_2 \cup A && \text{by the separation properties} \\
&= set_tree\ (inter\ t_1\ t_2) && \text{by (10.1), (10.5), IHs, } bst\ t_1, bst\ t_2, \text{(10.6)}
\end{aligned}$$

The correctness properties of *diff* follow the same pattern and their proofs are similar to the proofs of the *inter* properties. This concludes the generic join approach.

10.3 Joining Red-Black Trees

This section shows how to implement *join* efficiently on red-black trees. The basic idea is simple: descend along the spine of the higher of the two trees until reaching a subtree whose height is the same as the height of the lower tree. With suitable changes this works for other balanced trees as well [9]. The function definitions are shown in Figure 10.5. Function *join* calls *joinR* (descending along the right spine of l) if l is the higher tree, or calls *joinL* (descending along the left spine of r) if r is the higher tree, or returns $B\ l\ x\ r$

```
joinL :: 'a rbt ⇒ 'a ⇒ 'a rbt ⇒ 'a rbt
joinL l x r
= (if bh r ≤ bh l then R l x r
    else case r of
        ⟨l', (x', Red), r'⟩ ⇒ R (joinL l x l') x' r' |
        ⟨l', (x', Black), r'⟩ ⇒ baliL (joinL l x l') x' r')

joinR :: 'a rbt ⇒ 'a ⇒ 'a rbt ⇒ 'a rbt
joinR l x r
= (if bh l ≤ bh r then R l x r
    else case l of
        ⟨l', (x', Red), r'⟩ ⇒ R l' x' (joinR r' x r) |
        ⟨l', (x', Black), r'⟩ ⇒ baliR l' x' (joinR r' x r))

join :: 'a rbt ⇒ 'a ⇒ 'a rbt ⇒ 'a rbt
join l x r
= (if bh r < bh l then paint Black (joinR l x r)
    else if bh l < bh r then paint Black (joinL l x r) else B l x r)
```

Fig. 10.5. Function *join* on red-black trees

otherwise. The running time is linear in the black height (and thus logarithmic in the size) if we assume that the black height is stored in each node; our implementation of red-black trees would have to be augmented accordingly. Note that in *joinR* (and similarly in *joinL*) the comparison is not $bh\ l = bh\ r$ but $bh\ l \leq bh\ r$ to simplify the proofs.

10.3.1 Correctness

We need to prove that *join* on red-black trees (and a suitable *inv*) satisfies its specification in Figure 10.2. We start with properties of *joinL*; the properties of function *joinR* are completely symmetric. These are the automatically provable inductive properties:

$$
\begin{aligned}
&invc\ l \wedge invc\ r \wedge invh\ l \wedge invh\ r \wedge bh\ l \leq bh\ r \longrightarrow \\
&invc2\ (joinL\ l\ x\ r) \wedge \\
&(bh\ l \neq bh\ r \wedge color\ r = Black \longrightarrow invc\ (joinL\ l\ x\ r)) \wedge \\
&invh\ (joinL\ l\ x\ r) \wedge bh\ (joinL\ l\ x\ r) = bh\ r \\
&bh\ l \leq bh\ r \longrightarrow set_tree\ (joinL\ l\ x\ r) = set_tree\ l \cup \{x\} \cup set_tree\ r \\
&bst\ \langle l, (a, n), r\rangle \wedge bh\ l \leq bh\ r \longrightarrow bst\ (joinL\ l\ a\ r)
\end{aligned}
$$

Because *joinL* employs *baliL* from the chapter on red-black trees, the proof of the first property makes use of the *baliL* property shown in Section 8.2.1.

We define the invariant *inv* required for the specification in Figure 10.2 as follows:

$$inv\ t = (invc\ t \wedge inv\ h\ t)$$

Although weaker than *rbt*, it still guarantees logarithmic height (see Exercise 8.1). Note that *rbt* itself does not work because it does not satisfy property (10.4). The properties of *join* and *inv* are now easy consequences of the *joinL* (and *joinR*) properties shown above.

10.4 Exercises

Exercise 10.1. Define an alternative version *diff*1 of *diff* where in the third equation pattern matching is on t_1 and t_2 is *split*. Prove $bst\ t_1 \wedge bst\ t_2 \longrightarrow set_tree\ (diff1\ t_1\ t_2) = set_tree\ t_1 - set_tree\ t_2$.

Exercise 10.2. Following the general idea of the join function for red-black trees, define a join function for 2-3-trees. Start with two functions *joinL*, $joinR :: 'a\ tree23 \Rightarrow 'a \Rightarrow 'a\ tree23 \Rightarrow 'a\ upI$ and combine them into the overall join function:

$$join :: 'a\ tree23 \Rightarrow 'a \Rightarrow 'a\ tree23 \Rightarrow 'a\ tree23$$

Prove the following correctness properties:

$$complete\ l \wedge complete\ r \longrightarrow complete\ (join\ l\ x\ r)$$

$$complete\ l \wedge complete\ r \longrightarrow$$
$$inorder\ (join\ l\ x\ r) = inorder\ l\ @\ x\ \#\ inorder\ r$$

The corresponding (and needed) properties of *joinL* and *joinR* are slightly more involved.

Bibliographic Remarks

The join approach goes back to Adams [1]. Blelloch *et al.* [9] generalized the approach from weight-balanced trees to AVL trees, red-black trees and treaps. In particular they proved the $O(m \lg(\frac{n}{m} + 1))$ bound for the work (and an $O(\lg m \lg n)$ bound for the span).

11

Arrays via Braun Trees

Braun trees are a subclass of almost complete trees. In this chapter we explore their use as arrays and in Chapter 16 as priority queues.

11.1 Array

So far we have discussed sets (or maps) over some arbitrary linearly ordered type. Now we specialize that linearly ordered type to *nat* to model arrays. In principle we could model arrays as maps from a subset of natural numbers to the array elements. Because arrays are contiguous, it is more appropriate to model them as lists. The type *'a list* comes with two array-like operations (see Appendix A):

Indexing: $xs \mathbin{!} n$ is the nth element of the list xs.
Updating: $xs[n := x]$ is xs with the nth element replaced by x.

By convention, indexing starts with $n = 0$. If $n \geq \mathit{length}\ xs$ then $xs \mathbin{!} n$ and $xs[n := x]$ are underdefined: they are defined terms but we do not know what their value is.

Note that operationally, indexing and updating take time linear in the index, which may appear inappropriate for arrays. However, the type of lists is only an abstract model that specifies the desired functional behaviour of arrays but not their running time complexity.

The ADT of arrays is shown in Figure 11.1. Type *'ar* is the type of arrays, type *'a* the type of elements in the arrays. The abstraction function *list* abstracts arrays to lists. It would make perfect sense to include *list* in the interface as well. In fact, our implementation below comes with a (reasonably efficiently) executable definition of *list*.

The behaviour of *lookup*, *update*, *size* and *array* is specified in terms of their counterparts on lists and requires that the invariant is preserved.

ADT *Array* =

interface
lookup :: $'ar \Rightarrow nat \Rightarrow 'a$
update :: $nat \Rightarrow 'a \Rightarrow 'ar \Rightarrow 'ar$
len :: $'ar \Rightarrow nat$
array :: $'a\ list \Rightarrow 'ar$

abstraction *list* :: $'ar \Rightarrow 'a\ list$
invariant *invar* :: $'ar \Rightarrow bool$

specification

$invar\ ar \wedge n < len\ ar \longrightarrow lookup\ ar\ n = list\ ar\ !\ n$	(*lookup*)
$invar\ ar \wedge n < len\ ar \longrightarrow invar\ (update\ n\ x\ ar)$	(*update-inv*)
$invar\ ar \wedge n < len\ ar \longrightarrow list\ (update\ n\ x\ ar) = (list\ ar)[n := x]$	(*update*)
$invar\ ar \longrightarrow len\ ar = \lvert list\ ar\rvert$	(*len*)
$invar\ (array\ xs)$	(*array-inv*)
$list\ (array\ xs) = xs$	(*array*)

Fig. 11.1. ADT *Array*

What distinguishes the specifications of *lookup* and *update* from the standard schema (see Chapter 6) is that they carry a size precondition because the result of *lookup* and *update* is only specified if the index is less than the size of the array.

11.2 Braun Trees

One can implement arrays by any one of the many search trees presented in this book. Instead we take advantage of the fact that the keys are natural numbers and implement arrays by so-called **Braun trees** which are almost complete and thus have minimal height.

The basic idea is to index a node in a binary tree by the non-zero bit string that leads from the root to that node in the following fashion. Starting from the least significant bit and while we have not reached the leading 1 (which is ignored), we examine the bits one by one. If the current bit is 0, descend into the left child, otherwise into the right child. Instead of bit strings we use the natural numbers ≥ 1 that they represent. The Braun tree with nodes indexed by 1–15 is shown in Figure 11.2. The numbers are the indexes and not the elements stored in the nodes. For example, the index 14 is 0111 in binary (least significant bit first). If you follow the path left-right-right corresponding to 011 in Figure 11.2 you reach node 14.

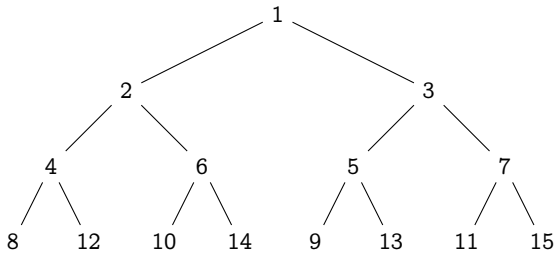


Fig. 11.2. Braun tree with nodes indexed by 1–15

A tree t is suitable for representing an array if the set of indexes of all its nodes is the interval $\{1..|t|\}$. The following tree is unsuitable because the node indexed by 2 is missing:



It turns out that the following invariant guarantees that a tree t contains exactly the nodes indexed by 1, ..., $|t|$:

$$\begin{aligned}&\mathit{braun} :: \mathit{'a\ tree} \Rightarrow \mathit{bool}\\ &\mathit{braun}\ \langle\rangle = \mathit{True}\\ &\mathit{braun}\ \langle l, _, r\rangle = ((|l| = |r| \vee |l| = |r| + 1) \wedge \mathit{braun}\ l \wedge \mathit{braun}\ r)\end{aligned}$$

The disjunction can alternatively be expresses as $|r| \leq |l| \leq |r| + 1$. We call a tree a **Braun tree** iff it satisfies predicate *braun*.

Although we do not need or prove this here, it is interesting to note that a tree that contains exactly the nodes indexed by 1, ..., $|t|$ is a Braun tree.

Let us now prove the earlier claim that Braun trees are almost complete. First, a lemma about the composition of almost complete trees:

Lemma 11.1.
$\mathit{acomplete}\ l \wedge \mathit{acomplete}\ r \wedge |l| = |r| + 1 \longrightarrow \mathit{acomplete}\ \langle l, x, r\rangle$

Proof. Using Lemmas 4.7 and 4.8 and the assumptions we obtain

$$h\ \langle l, x, r\rangle = \lceil \lg\ (|r|_1 + 1)\rceil + 1 \qquad (*)$$
$$mh\ \langle l, x, r\rangle = \lfloor \lg\ |r|_1\rfloor + 1 \qquad (**)$$

Because $1 \leq |r|_1$ there is an i such that $2^i \leq |r|_1 < 2^{i+1}$ and thus $2^i < |r|_1 + 1 \leq 2^{i+1}$. This implies $i = \lfloor \lg\ |r|_1\rfloor$ and $i + 1 = \lceil \lg\ (|r|_1 + 1)\rceil$. Together with $(*)$ and $(**)$ this implies $\mathit{acomplete}\ \langle l, x, r\rangle$. □

Now we can show that all Braun trees are almost complete:

Lemma 11.2. *braun t* $\longrightarrow$ *acomplete t*

Thus we know that Braun trees have optimal height (Lemma 4.6) and can even quantify it (Lemma 4.7).

Proof. The proof is by induction. We focus on the induction step where $t = \langle l, x, r \rangle$. By assumption we have *acomplete l* and *acomplete r*. Because of *braun t* we can distinguish two cases. First assume $|l| = |r| + 1$. The claim *acomplete t* follows immediately from the previous lemma. Now assume $|l| = |r|$. By definition, there are four cases to consider when proving *acomplete t*. By symmetry it suffices to consider only two of them. If $h\ l \leq h\ r$ and $mh\ r < mh\ l$ then *acomplete t* reduces to *acomplete r*, which is true by assumption. Now assume $h\ l \leq h\ r$ and $mh\ l \leq mh\ r$. Because $|l| = |r|$, the fact that the height of an almost complete tree is determined uniquely by its size (Lemma 4.7) implies $h\ l = h\ r$ and thus *acomplete t* reduces to *acomplete l*, which is again true by assumption. □

Note that the proof does not rely on the fact that it is the left child that is potentially one bigger than the right one; it merely requires that the difference in size between two siblings is at most 1.

11.3 Arrays via Braun Trees

In this section we implement arrays by means of Braun trees and verify correctness and complexity. We start by defining array-like functions on Braun trees. After the above explanation of Braun trees the following lookup function will not come as a surprise:

```
lookup1 :: 'a tree ⇒ nat ⇒ 'a
lookup1 ⟨l, x, r⟩ n
= (if n = 1 then x else lookup1 (if even n then l else r) (n div 2))
```

The least significant bit is the parity of the index and we advance to the next bit by *div* 2. The function is called *lookup*1 rather than *lookup* to emphasize that it expects the index to be at least 1. This simplifies the implementation via Braun trees but is in contrast to the *Array* interface where by convention indexing starts with 0.

Function *update*1 descends in the very same manner but also performs an update when reaching 1:

```
update1 :: nat ⇒ 'a ⇒ 'a tree ⇒ 'a tree
update1 _ x ⟨⟩ = ⟨⟨⟩, x, ⟨⟩⟩
update1 n x ⟨l, a, r⟩
= (if n = 1 then ⟨l, x, r⟩
    else if even n then ⟨update1 (n div 2) x l, a, r⟩
         else ⟨l, a, update1 (n div 2) x r⟩)
```

The second equation updates existing entries in case $n = 1$. The first equation, however, creates a new entry and thus supports extending the tree. That is, *update*1 $(|t| + 1)$ x t extends the tree with a new node x at index $|t| + 1$. Function *adds* iterates this process (again expecting $|t| + 1$ as the index) and thus adds a whole list of elements:

```
adds :: 'a list ⇒ nat ⇒ 'a tree ⇒ 'a tree
adds [] _ t = t
adds (x # xs) n t = adds xs (n + 1) (update1 (n + 1) x t)
```

The implementation of the *Array* interface in Figure 11.3 is just a thin wrapper around the corresponding functions on Braun trees. An array is represented as a pair of a Braun tree and its size. Note that although *update*1 can extend the tree, the specification and implementation of the array *update* function does not support that: n is expected to be below the length of the array. Flexible arrays are specified and implemented in Section 11.4.

```
lookup (t, _) n    = lookup1 t (n + 1)
update n x (t, m) = (update1 (n + 1) x t, m)
len (t, m)         = m
array xs           = (adds xs 0 ⟨⟩, |xs|)
```

Fig. 11.3. Array implementation via Braun trees

11.3.1 Functional Correctness

The invariant on arrays is obvious:

$$invar\ (t,\ l) = (braun\ t \wedge l = |t|)$$

The abstraction function *list* could be defined in the following intuitive way, where $[m..<n]$ is the list of natural numbers from m up to but excluding n (see Appendix A):

$$list\ t = map\ (lookup1\ t)\ [1..<|t| + 1]$$

Instead we define *list* recursively and derive the above equation later on

$$\begin{array}{l} list :: 'a\ tree \Rightarrow 'a\ list \\ list\ \langle\rangle = [] \\ list\ \langle l,\ x,\ r\rangle = x\ \#\ splice\ (list\ l)\ (list\ r) \end{array}$$

This definition is best explained by looking at Figure 11.2. The subtrees with root 2 and 3 will be mapped to the lists [2, 4, 6, 8, 10, 12, 14] and [1, 3, 5, 7, 9, 11, 13, 15]. The obvious way to combine these two lists into [1, 2, 3, ..., 15] is to splice them:

$$\begin{array}{l} splice :: 'a\ list \Rightarrow 'a\ list \Rightarrow 'a\ list \\ splice\ []\ ys = ys \\ splice\ (x\ \#\ xs)\ ys = x\ \#\ splice\ ys\ xs \end{array}$$

Note that because of this reasonably efficient ($O(n \lg n)$, see Section 11.3.2) implementation of *list* we can also regard *list* as part of the interface of arrays.

Before we embark on the actual proofs we state a helpful arithmetic truth that is frequently used implicitly below:

$$\begin{array}{l} braun\ \langle l,\ x,\ r\rangle \wedge n \in \{1..|\langle l,\ x,\ r\rangle|\} \wedge 1 < n \longrightarrow \\ (odd\ n \longrightarrow n\ div\ 2 \in \{1..|r|\}) \wedge (even\ n \longrightarrow n\ div\ 2 \in \{1..|l|\}) \end{array}$$

where $\{m..n\} = \{k \mid m \leq k \wedge k \leq m\}$.

We will now verify that the implementation in Figure 11.3 of the *Array* interface in Figure 11.1 satisfies the given specification.

We start with property (*len*), the correctness of function *len*. Because of the invariant, (*len*) follows directly from

$$|list\ t| = |t|$$

which is proved by induction. We will also use this property implicitly in many proofs below.

The following proposition implies the correctness property (*lookup*):

$$braun\ t \wedge i < |t| \longrightarrow list\ t\ !\ i = lookup1\ t\ (i + 1) \tag{11.1}$$

The proof is by induction and uses the following proposition that is also proved by induction:

$$n < |xs| + |ys| \wedge |ys| \leq |xs| \wedge |xs| \leq |ys| + 1 \longrightarrow$$
$$splice\ xs\ ys\ !\ n = (\textbf{if}\ even\ n\ \textbf{then}\ xs\ \textbf{else}\ ys)\ !\ (n\ div\ 2)$$

As a corollary to (11.1) we obtain that function *list* can indeed be expressed via *lookup*1:

$$braun\ t \longrightarrow list\ t = map\ (lookup1\ t)\ [1..<|t| + 1] \tag{11.2}$$

It follows by **list extensionality**:

$$xs = ys \longleftrightarrow |xs| = |ys| \wedge (\forall\, i{<}|xs|.\ xs\ !\ i = ys\ !\ i)$$

Let us now verify *update* as implemented via *update*1. The following two preservation properties (proved by induction) prove (*update-inv*):

$$braun\ t \wedge n \in \{1..|t|\} \longrightarrow |update1\ n\ x\ t| = |t|$$
$$braun\ t \wedge n \in \{1..|t|\} \longrightarrow braun\ (update1\ n\ x\ t)$$

The following property relating *lookup*1 and *update*1 is again proved by induction:

$$braun\ t \wedge n \in \{1..|t|\} \longrightarrow$$
$$lookup1\ (update1\ n\ x\ t)\ m = (\textbf{if}\ n = m\ \textbf{then}\ x\ \textbf{else}\ lookup1\ t\ m)$$

The last three properties together with (11.2) and list extensionality prove the following proposition, which implies (*update*):

$$braun\ t \wedge n \in \{1..|t|\} \longrightarrow list\ (update1\ n\ x\ t) = (list\ t)[n - 1 := x]$$

Finally we turn to the constructor *array*. It is implemented in terms of *adds* and *update*1. Their correctness is captured by the following properties whose inductive proofs build on each other:

$$braun\ t \longrightarrow |update1\ (|t| + 1)\ x\ t| = |t| + 1 \tag{11.3}$$
$$braun\ t \longrightarrow braun\ (update1\ (|t| + 1)\ x\ t) \tag{11.4}$$
$$braun\ t \longrightarrow list\ (update1\ (|t| + 1)\ x\ t) = list\ t\ @\ [x] \tag{11.5}$$
$$braun\ t \longrightarrow |adds\ xs\ |t|\ t| = |t| + |xs| \wedge braun\ (adds\ xs\ |t|\ t)$$
$$braun\ t \longrightarrow list\ (adds\ xs\ |t|\ t) = list\ t\ @\ xs$$

The last two properties imply the remaining proof obligations (*array-inv*) and (*array*). The proof of (11.5) requires the following two properties of *splice* which are proved by simultaneous induction:

$$|ys| \leq |xs| \longrightarrow splice\ (xs\ @\ [x])\ ys = splice\ xs\ ys\ @\ [x]$$
$$|xs| \leq |ys| + 1 \longrightarrow splice\ xs\ (ys\ @\ [y]) = splice\ xs\ ys\ @\ [y]$$

11.3.2 Running Time Analysis

The running time of *lookup* and *update* is obviously logarithmic because of the logarithmic height of Braun trees. We sketch why *list* and *array* both have running time $O(n \lg n)$. Linear time versions are presented in Section 11.5.

Function *list* is similar to bottom-up merge sort and *splice* is similar to *merge*. We focus on *splice* because it performs almost all the work. Consider calling *list* on a complete tree of height h. At each level k (starting with 0 for the root) of the tree, *splice* is called 2^k times with lists of size (almost) 2^{h-k-1}. The running time of *splice* with lists of the same length is proportional to the size of the lists. Thus the running time at each level is $O(2^k 2^{h-k-1}) = O(2^{h-1}) = O(2^h)$. Thus all the splices together require time $O(h\,2^h)$. Because complete trees have size $n = 2^h$, the bound $O(n \lg n)$ follows.

Function *array* is implemented via *adds* and thus via repeated calls of *update*1. At the beginning of Section 7.3 we show that because *update*1 has logarithmic complexity, calling it n times on a growing tree starting with a leaf takes time $\Theta(n \lg n)$.

11.4 Flexible Arrays

Flexible arrays can be grown and shrunk at either end. Figure 11.4 shows the specification of all four operations. (For *tl* and *butlast* see Appendix A.) *Array_Flex* extends the basis specification *Array* in Figure 11.1.

ADT *Array_Flex* = *Array* +

interface

add_lo :: $'a \Rightarrow {}'ar \Rightarrow {}'ar$
del_lo :: $'ar \Rightarrow {}'ar$
add_hi :: $'a \Rightarrow {}'ar \Rightarrow {}'ar$
del_hi :: $'ar \Rightarrow {}'ar$

specification

invar ar $\longrightarrow$ *invar* (*add_lo a ar*)	(*add_lo-inv*)
invar ar $\longrightarrow$ *list* (*add_lo a ar*) = *a* # *list ar*	(*add_lo*)
invar ar $\longrightarrow$ *invar* (*del_lo ar*)	(*del_lo-inv*)
invar ar $\longrightarrow$ *list* (*del_lo ar*) = *tl* (*list ar*)	(*del_lo*)
invar ar $\longrightarrow$ *invar* (*add_hi a ar*)	(*add_hi-inv*)
invar ar $\longrightarrow$ *list* (*add_hi a ar*) = *list ar* @ [*a*]	(*add_hi*)
invar ar $\longrightarrow$ *invar* (*del_hi ar*)	(*del_hi-inv*)
invar ar $\longrightarrow$ *list* (*del_hi ar*) = *butlast* (*list ar*)	(*del_hi*)

Fig. 11.4. ADT *Array_Flex*

Below we first implement the *Array_Flex* functions on Braun trees. In a final step an implementation of *Array_Flex* on (tree,size) pairs is derived.

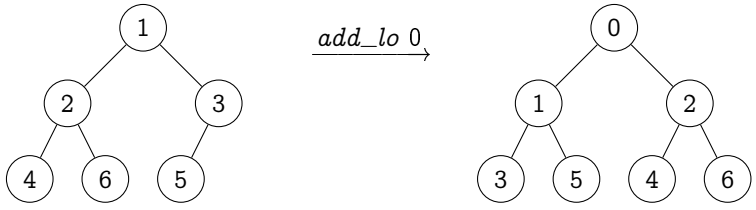
We have already seen that *update*1 adds an element at the high end. The inverse operation *del_hi* removes the high end, assuming that the given index is the size of the tree:

```
del_hi :: nat ⇒ 'a tree ⇒ 'a tree
del_hi _ ⟨⟩ = ⟨⟩
del_hi n ⟨l, x, r⟩
= (if n = 1 then ⟨⟩
   else if even n then ⟨del_hi (n div 2) l, x, r⟩
        else ⟨l, x, del_hi (n div 2) r⟩)
```

This was easy but extending an array at the low end seems hard because one has to shift the existing entries. However, Braun trees support a logarithmic implementation:

```
add_lo :: 'a ⇒ 'a tree ⇒ 'a tree
add_lo x ⟨⟩ = ⟨⟨⟩, x, ⟨⟩⟩
add_lo x ⟨l, a, r⟩ = ⟨add_lo a r, x, l⟩
```

The intended functionality is *list* (*add_lo x t*) = *x* # *list t*. Function *add_lo* installs the new element *x* at the root of the tree. Because *add_lo* needs to shift the indices of the elements already in the tree, the left child (indices 2, 4, ...) becomes the new right child (indices 3, 5, ...). The old right child becomes the new left child with the old root *a* added in at index 2 and the remaining elements at indices 4, 6, ... In the following example, *add_lo* 0 transforms the left tree into the right one. The numbers in the nodes are the actual elements, not their indices.



Function *del_lo* simply reverses *add_lo* by removing the root and merging the children:

```
del_lo :: 'a tree ⇒ 'a tree
del_lo ⟨⟩ = ⟨⟩
del_lo ⟨l, _, r⟩ = merge l r
```

```
merge :: 'a tree ⇒ 'a tree ⇒ 'a tree
merge ⟨⟩ r = r
merge ⟨l, a, r⟩ rr = ⟨rr, a, merge l r⟩
```

Figure 11.5 shows the obvious implementation of the functions in the *Array_Flex* specification from Figure 11.4 (on the left-hand side) with the help of the corresponding Braun tree operations (on the right-hand side). It is an extension of the basic array implementation from Figure 11.3. All *Array_Flex* functions have logarithmic time complexity because the corresponding Braun tree functions do because they descend along one branch of the tree.

```
add_lo x (t, l) = (add_lo x t, l + 1)
del_lo (t, l)   = (del_lo t, l − 1)
add_hi x (t, l) = (update1 (l + 1) x t, l + 1)
del_hi (t, l)   = (del_hi l t, l − 1)
```

Fig. 11.5. Flexible array implementation via Braun trees

11.4.1 Functional Correctness

We now have to prove the properties in Figure 11.4. We have already dealt with *update*1 and thus *add_hi* above. Properties (*add_hi-inv*) and (*add_hi*) follow from (11.3), (11.4) and (11.5) stated earlier.

Correctness of *del_hi* on Braun trees is captured by the following two properties proved by induction:

$$braun\ t \longrightarrow braun\ (del_hi\ |t|\ t)$$
$$braun\ t \longrightarrow list\ (del_hi\ |t|\ t) = butlast\ (list\ t) \tag{11.6}$$

They imply (*del_hi-inv*) and (*del_hi*). The proof of (11.6) requires the simple fact $list\ t = [] \longleftrightarrow t = \langle\rangle$ and the following property of *splice* which is proved by induction:

$butlast\ (splice\ xs\ ys)$
$= (\textbf{if}\ |ys| < |xs|\ \textbf{then}\ splice\ (butlast\ xs)\ ys\ \textbf{else}\ splice\ xs\ (butlast\ ys))$

Correctness of *add_lo* on Braun trees is captured by the following two properties proved by induction:

$braun\ t \longrightarrow braun\ (add_lo\ x\ t)$

$braun\ t \longrightarrow list\ (add_lo\ a\ t) = a\ \#\ list\ t$

Properties (*add_lo-inv*) and (*add_lo*) follow directly from them.

Finally we turn to *del_lo*. Inductions (for *merge*) and case analyses (for *del_lo*) yield the following properties:

$braun\ \langle l,\ x,\ r\rangle \longrightarrow braun\ (merge\ l\ r)$

$braun\ \langle l,\ x,\ r\rangle \longrightarrow list\ (merge\ l\ r) = splice\ (list\ l)\ (list\ r)$

$braun\ t \longrightarrow braun\ (del_lo\ t)$

$braun\ t \longrightarrow list\ (del_lo\ t) = tl\ (list\ t)$

The last two properties imply (*del_lo-inv*) and (*del_lo*).

11.5 Bigger, Better, Faster, More!

In this section we meet efficient versions of some old and new functions on Braun trees. The implementation of the corresponding array operations is trivial and is not discussed.

11.5.1 Fast Size of Braun Trees

The size of a Braun tree can be computed without having to traverse the entire tree:

```
size_fast :: 'a tree ⇒ nat
size_fast ⟨⟩ = 0
size_fast ⟨l, _, r⟩ = (let n = size_fast r in 1 + 2 · n + diff l n)

diff :: 'a tree ⇒ nat ⇒ nat
diff ⟨⟩ _ = 0
diff ⟨l, _, r⟩ n
= (if n = 0 then 1
   else if even n then diff r (n div 2 − 1) else diff l (n div 2))
```

Function *size_fast* descends down the right spine, computes the size of a *Node* as if both children were the same size $(1 + 2 \cdot n)$, but adds *diff l n* to compensate for bigger left children. Correctness of *size_fast*

Lemma 11.3. $braun\ t \longrightarrow size_fast\ t = |t|$

follows from this property of *diff*:

$$braun\ t \wedge |t| \in \{n,\ n+1\} \longrightarrow diff\ t\ n = |t| - n$$

The running time of *size_fast* is quadratic in the height of the tree (see Exercise 11.3).

11.5.2 Initializing a Braun Tree with a Fixed Value

Above we only considered the construction of a Braun tree from a list. Alternatively one may want to create a tree (array) where all elements are initialized to the same value. Of course one can call *update*1 n times, but one can also build the tree directly:

```
braun_of_naive x n
= (if n = 0 then ⟨⟩
   else let m = (n − 1) div 2
        in if odd n
           then ⟨braun_of_naive x m, x, braun_of_naive x m⟩
           else ⟨braun_of_naive x (m + 1), x,
                 braun_of_naive x m⟩)
```

This solution also has time complexity $O(n \lg n)$ but it can clearly be improved by sharing identical recursive calls. Function *braun2_of* shares as much as possible by producing trees of size n and $n + 1$ in parallel:

```
braun2_of :: 'a ⇒ nat ⇒ 'a tree × 'a tree
braun2_of x n
= (if n = 0 then (⟨⟩, ⟨⟨⟩, x, ⟨⟩⟩)
   else let (s, t) = braun2_of x ((n − 1) div 2)
        in if odd n then (⟨s, x, s⟩, ⟨t, x, s⟩) else (⟨t, x, s⟩, ⟨t, x, t⟩))

braun_of :: 'a ⇒ nat ⇒ 'a tree
braun_of x n = fst (braun2_of x n)
```

The running time is clearly logarithmic.

The correctness properties (see Appendix A for *replicate*)

$$list\ (braun_of\ x\ n) = replicate\ n\ x$$
$$braun\ (braun_of\ x\ n)$$

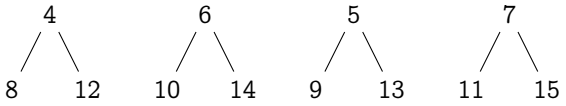
are corollaries of the more general statements

$$braun2_of\ x\ n = (s,\ t) \longrightarrow$$
$$list\ s = replicate\ n\ x \wedge list\ t = replicate\ (n + 1)\ x$$
$$braun2_of\ x\ n = (s,\ t) \longrightarrow$$
$$|s| = n \wedge |t| = n + 1 \wedge braun\ s \wedge braun\ t$$

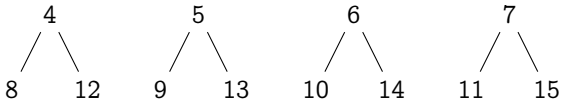
which can both be proved by induction.

11.5.3 Converting a List into a Braun Tree

We improve on function *adds* from Section 11.3 that has running time $\Theta(n \lg n)$ by developing a linear-time function. Given a list of elements $[1, 2, \ldots]$, we can subdivide it into sublists $[1]$, $[2, 3]$, $[4, \ldots, 7]$, $\ldots$ such that the kth sublist contains the elements of level k of the corresponding Braun tree. This is simply because on each level we have the entries whose index has $k+1$ bits. Thus we need to process the input list in chunks of size 2^k to produce the trees on level k. But we also need to get the order right. To understand how that works, consider the last two levels of the tree in Figure 11.2:



If we rearrange them in increasing order of the root labels



the following pattern emerges: the left subtrees are labeled $[8, \ldots, 11]$, the right subtrees $[12, \ldots, 15]$. Call t_i the tree with root label i. The correct order of subtrees, i.e. t_4, t_6, t_5, t_7, is restored when the three lists $[t_4,\ t_5]$, $[2,\ 3]$ (the labels above) and $[t_6,\ t_7]$ are combined into new trees by going through them simultaneously from left to right, yielding $[\langle t_4,\ 2,\ t_6\rangle,\ \langle t_5,\ 3,\ t_7\rangle]$, the level above.

Abstracting from this example we arrive at the following code. Loosely speaking, *brauns k xs* produces the Braun trees on level k.

```
brauns :: nat ⇒ 'a list ⇒ 'a tree list
brauns k xs
= (if xs = [] then []
   else let ys = take 2^k xs;
            zs = drop 2^k xs;
            ts = brauns (k + 1) zs
        in nodes ts ys (drop 2^k ts))
```

Function *brauns* chops off a chunk *ys* of size 2^k from the input list and recursively converts the remainder of the list into a list *ts* of (at most) 2^{k+1} trees. This list is (conceptually) split into *take* 2^k *ts* and *drop* 2^k *ts* which are combined with *ys* by function *nodes* that traverses its three argument lists simultaneously. As a local optimization, we pass all of *ts* rather than just *take* 2^k *ts* to *nodes*.

```
nodes :: 'a tree list ⇒ 'a list ⇒ 'a tree list ⇒ 'a tree list
nodes (l # ls) (x # xs) (r # rs) = ⟨l, x, r⟩ # nodes ls xs rs
nodes (l # ls) (x # xs) [] = ⟨l, x, ⟨⟩⟩ # nodes ls xs []
nodes [] (x # xs) (r # rs) = ⟨⟨⟩, x, r⟩ # nodes [] xs rs
nodes [] (x # xs) [] = ⟨⟨⟩, x, ⟨⟩⟩ # nodes [] xs []
nodes _ [] _ = []
```

Because the input list may not have exactly $2^n - 1$ elements, some of the chunks of elements and trees may be shorter than 2^k. To compensate for that, function *nodes* implicitly pads lists of trees at the end with leaves. This padding is the purpose of equations two to four.

The top-level function for turning a list into a tree simply extracts the first (and only) element from the list computed by *brauns* 0:

```
brauns1 :: 'a list ⇒ 'a tree
brauns1 xs = (if xs = [] then ⟨⟩ else brauns 0 xs ! 0)
```

Functional Correctness

The key correctness lemma below expresses a property of Braun trees: the subtrees on level k consist of all elements of the input list *xs* that are 2^k elements apart, starting from some offset. To state this concisely we define

```
take_nths :: nat ⇒ nat ⇒ 'a list ⇒ 'a list
take_nths _ _ [] = []
take_nths i k (x # xs)
= (if i = 0 then x # take_nths (2^k − 1) k xs
   else take_nths (i − 1) k xs)
```

The result of *take_nths i k xs* is every 2^k-th element in *drop i xs*.

A number of simple properties follow by easy inductions:

$$take_nths\ i\ k\ (drop\ j\ xs) = take_nths\ (i + j)\ k\ xs \tag{11.7}$$

$$take_nths\ 0\ 0\ xs = xs \tag{11.8}$$

$$splice\ (take_nths\ 0\ 1\ xs)\ (take_nths\ 1\ 1\ xs) = xs \tag{11.9}$$

$$take_nths\ i\ m\ (take_nths\ j\ n\ xs) = take_nths\ (i \cdot 2^n + j)\ (m + n)\ xs \tag{11.10}$$

$$take_nths\ i\ k\ xs = [] \longleftrightarrow |xs| \leq i \tag{11.11}$$

$$i < |xs| \longrightarrow hd\ (take_nths\ i\ k\ xs) = xs\ !\ i \tag{11.12}$$

$$|xs| = |ys| \vee |xs| = |ys| + 1 \longrightarrow take_nths\ 0\ 1\ (splice\ xs\ ys) = xs \wedge take_nths\ 1\ 1\ (splice\ xs\ ys) = ys \tag{11.13}$$

$$|take_nths\ 0\ 1\ xs| = |take_nths\ 1\ 1\ xs| \vee |take_nths\ 0\ 1\ xs| = |take_nths\ 1\ 1\ xs| + 1 \tag{11.14}$$

We also introduce a predicate relating a tree to a list:

```
braun_list :: 'a tree ⇒ 'a list ⇒ bool
braun_list ⟨⟩ xs = (xs = [])
braun_list ⟨l, x, r⟩ xs
= (xs ≠ [] ∧ x = hd xs ∧
    braun_list l (take_nths 1 1 xs) ∧
    braun_list r (take_nths 2 1 xs))
```

This definition may look a bit mysterious at first but it satisfies a simple specification: *braun_list t xs* $\longleftrightarrow$ *braun t* $\wedge$ *xs = list t* (see below). The idea of the above definition is that instead of relating $\langle l, x, r\rangle$ to *xs* via *splice* we invert the process and relate *l* and *r* to the even and odd numbered elements of *drop 1 xs*.

Lemma 11.4. *braun_list t xs* $\longleftrightarrow$ *braun t* $\wedge$ *xs = list t*

Proof. The proof is by induction on t. The base case is trivial. In the induction step the key properties are (11.14) to prove $braun\ t$ and (11.9) and (11.13) to prove $xs = list\ t$. □

The correctness proof of *brauns* rests on a few simple inductive properties:

$$|nodes\ ls\ xs\ rs| = |xs| \tag{11.15}$$

$$\begin{aligned}&i < |xs| \longrightarrow\\ &nodes\ ls\ xs\ rs\ !\ i\\ &= \langle \textbf{if } i < |ls| \textbf{ then } ls\ !\ i \textbf{ else } \langle\rangle,\ xs\ !\ i,\\ &\quad \textbf{if } i < |rs| \textbf{ then } rs\ !\ i \textbf{ else } \langle\rangle\rangle\end{aligned} \tag{11.16}$$

$$|brauns\ k\ xs| = min\ |xs|\ 2^k \tag{11.17}$$

The main theorem expresses the following correctness property of the elements of $brauns\ k\ xs$: every tree $brauns\ k\ xs\ !\ i$ is a Braun tree and its list of elements is $take_nths\ i\ k\ xs$:

Theorem 11.5. $i < min\ |xs|\ 2^k \longrightarrow$
$braun_list\ (brauns\ k\ xs\ !\ i)\ (take_nths\ i\ k\ xs)$

Proof. The proof is by induction on the length of xs. Assume $i < min\ |xs|\ 2^k$, which implies $xs \neq []$. Let $zs = drop\ 2^k\ xs$. Thus $|zs| < |xs|$ and therefore the IH applies to zs and yields the property

$$\begin{aligned}&\forall i\ j.\ j = i + 2^k \wedge i < min\ |zs|\ 2^{k+1} \longrightarrow\\ &\qquad braun_list\ (ts\ !\ i)\ (take_nths\ j\ (k+1)\ xs)\end{aligned} \tag{$*$}$$

where $ts = brauns\ (k+1)\ zs$. Let $ts' = drop\ 2^k\ ts$. Below we examine $nodes\ ts\ _\ ts'\ !\ i$ with the help of (11.16). Thus there are four similar cases of which we only discuss one representative one: assume $i < |ts|$ and $i \geq |ts'|$.

$$\begin{aligned}&braun_list\ (brauns\ k\ xs\ !\ i)\ (take_nths\ i\ k\ xs)\\ &\longleftrightarrow braun_list\ (nodes\ ts\ (take\ 2^k\ xs)\ ts'\ !\ i)\ (take_nths\ i\ k\ xs)\\ &\longleftrightarrow braun_list\ (ts\ !\ i)\ (take_nths\ (2^k + i)\ (k+1)\ xs) \wedge\\ &\quad\ braun_list\ \langle\rangle\ (take_nths\ (2^{k+1} + i)\ (k+1)\ xs)\\ &\qquad\qquad \text{by (11.16), (11.10), (11.11), (11.12) and assumptions}\\ &\longleftrightarrow True \qquad\qquad \text{by ($*$), (11.11), (11.17) and assumptions}\end{aligned}$$

□

Setting $i = k = 0$ in this theorem we obtain the correctness of $brauns1$ using Lemma 11.4 and (11.8):

Corollary 11.6. $braun\ (brauns1\ xs) \wedge list\ (brauns1\ xs) = xs$

Running Time Analysis

We focus on function *brauns*. In the step from *brauns* to T_{brauns} we simplify matters a little bit: we count only the expensive operations that traverse lists and ignore the other small additive constants. The time to evaluate *take n xs* and *drop n xs* is linear in *min n* $|xs|$ and we simply use *min n* $|xs|$. Evaluating *nodes ls xs rs* takes time linear in $|xs|$ and $|take\ n\ xs| = min\ n\ |xs|$. As a result we obtain the following definition of T_{brauns}:

$$
\begin{array}{l}
T_{brauns} :: nat \Rightarrow 'a\ list \Rightarrow nat \\
T_{brauns}\ k\ xs \\
= (\textbf{if}\ xs = []\ \textbf{then}\ 0 \\
\quad \textbf{else let}\ ys = take\ 2^k\ xs;\ zs = drop\ 2^k\ xs;\ ts = brauns\ (k+1)\ zs \\
\qquad \textbf{in}\ 4 \cdot min\ 2^k\ |xs| + T_{brauns}\ (k+1)\ zs)
\end{array}
$$

It is easy to prove that T_{brauns} is linear:

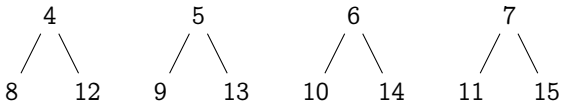
Lemma 11.7. $T_{brauns}\ k\ xs = 4 \cdot |xs|$

Proof. The proof is by induction on the length of *xs*. If $xs = []$ the claim is trivial. Now assume $xs \neq []$ and let $zs = drop\ 2^k\ xs$.

$$
\begin{array}{ll}
T_{brauns}\ k\ xs = T_{brauns}\ (k+1)\ zs + 4 \cdot min\ 2^k\ |xs| & \\
= 4 \cdot |zs| + 4 \cdot min\ 2^k\ |xs| & \text{by IH} \\
= 4 \cdot (|xs| - 2^k) + 4 \cdot min\ 2^k\ |xs| = 4 \cdot |xs| & \square
\end{array}
$$

11.5.4 Converting a Braun Tree into a List

We improve on function *list* that has running time $O(n \lg n)$ by developing a linear-time version. Imagine that we want to invert the computation of *brauns*1 and thus of *brauns*. Thus it is natural to convert not merely a single tree but a list of trees. Looking once more at the reordered list of subtrees



the following strategy strongly suggests itself: first the roots, then the left children, then the right children. The recursive application of this strategy also takes care of the required reordering of the subtrees. Of course we have to ignore any leaves we encounter. This is the resulting function:

```
list_fast_rec :: 'a tree list ⇒ 'a list
list_fast_rec ts
= (let us = filter (λt. t ≠ ⟨⟩) ts
   in if us = [] then []
      else map value us @ list_fast_rec (map left us @ map right us))

value ⟨l, x, r⟩ = x
left ⟨l, x, r⟩ = l
right ⟨l, x, r⟩ = r
```

Termination of *list_fast_rec* is almost obvious because *left* and *right* remove the top node of a tree. Thus *size* seems the right measure. But if $ts = [\langle\rangle]$, the measure is 0 but it still leads to a recursive call (with argument []). This problem can be avoided with the measure function $\varphi = sum_list \circ map\ f$ where $f = (\lambda t.\ 2 \cdot |t| + 1)$. Assume $ts \neq []$ and let $us = filter\ (\lambda t.\ t \neq \langle\rangle)\ ts$. We need to show that $\varphi\ (map\ left\ us\ @\ map\ right\ us) < \varphi\ ts$. Take some t in ts. If $t = \langle\rangle$, $f\ t = 1$ but t is no longer in us, i.e. the measure decreases by 1. If $t = \langle l, x, r\rangle$ then $f\ t = 2 \cdot |l| + 2 \cdot |r| + 3$ but $f\ (left\ t) + f\ (right\ t) = 2 \cdot |l| + 2 \cdot |r| + 2$ and thus the measure also decreases by 1. Because $ts \neq []$ this proves $\varphi\ (map\ left\ us\ @\ map\ right\ us) < \varphi\ ts$. We do not show the technical details.

Finally, the top level function to extract a list from a single tree:

```
list_fast :: 'a tree ⇒ 'a list
list_fast t = list_fast_rec [t]
```

From *list_fast* one can easily derive an efficient fold function on Braun trees that processes the elements in the tree in the order of their indexes.

Functional Correctness

We want to prove correctness of *list_fast*: $list_fast\ t = list\ t$ if $braun\ t$. A direct proof of $list_fast_rec\ [t] = list\ t$ will fail and we need to generalize this statement to all lists of length 2^k. Reusing the infrastructure from the previous subsection this can be expressed as follows:

Theorem 11.8. $|ts| = 2^k \wedge (\forall i{<}2^k.\ braun_list\ (ts\ !\ i)\ (take_nths\ i\ k\ xs)) \longrightarrow list_fast_rec\ ts = xs$

Proof. The proof is by induction on the length of xs. Assume the two premises. There are two cases.

First assume $|xs| < 2^k$. Then

$$ts = map\ (\lambda x.\ \langle\langle\rangle,\ x,\ \langle\rangle\rangle)\ xs\ @\ replicate\ n\ \langle\rangle \qquad (*)$$

where $n = |ts| - |xs|$. This can be proved pointwise. Take some $i < 2^k$. If $i < |xs|$ then $take_nths\ i\ k\ xs = take\ 1\ (drop\ i\ xs)$ (which can be proved by induction on xs). By definition of $braun_list$ it follows that $t\ !\ i = \langle l,\ xs\ !\ i,\ r\rangle$ for some l and r such that $braun_list\ l\ []$ and $braun_list\ l\ []$ and thus $l = r = \langle\rangle$, i.e. $t\ !\ i = \langle\langle\rangle,\ xs\ !\ i,\ \langle\rangle\rangle$. If $\neg\ i < |xs|$ then $take_nths\ i\ k\ xs = []$ by (11.11) and thus $braun_list\ (ts\ !\ i)\ []$ by the second premise and thus $ts\ !\ i = \langle\rangle$ by definition of $braun_list$. This concludes the proof of $(*)$. The desired $list_fast_rec\ ts = xs$ follows easily by definition of $list_fast_rec$.

Now assume $\neg\ |xs| < 2^k$. Then for all $i < 2^k$

$$\begin{aligned}&ts\ !\ i \neq \langle\rangle \wedge value\ (ts\ !\ i) = xs\ !\ i \wedge\\ &braun_list\ (left\ (ts\ !\ i))\ (take_nths\ (i + 2^k)\ (k + 1)\ xs) \wedge\\ &braun_list\ (right\ (ts\ !\ i))\ (take_nths\ (i + 2 \cdot 2^k)\ (k + 1)\ xs)\end{aligned}$$

follows from the second premise with the help of (11.10), (11.11) and (11.12). We obtain two consequences:

$$map\ value\ ts = take\ 2^k\ xs$$

$$list_fast_rec\ (map\ left\ ts\ @\ map\ right\ ts) = drop\ 2^k\ xs$$

The first consequence follows by pointwise reasoning, the second consequence with the help of the IH and (11.7). From these two consequences the desired conclusion $list_fast_rec\ ts = xs$ follows by definition of $list_fast_rec$. □

Running Time Analysis

We focus on $list_fast_rec$. In the step from $list_fast_rec$ to $T_{list_fast_rec}$ we simplify matters a little bit: we count only the expensive operations that traverse lists and ignore the other small additive constants. The time to evaluate $map\ value\ ts$, $map\ left\ ts$, $map\ right\ ts$, $filter\ (\lambda t.\ t \neq \langle\rangle)\ ts$ and $ts\ @\ _$ is linear in $|ts|$ and we simply use $|ts|$. As a result we obtain the following definition of $T_{list_fast_rec}$:

$$\begin{aligned}&T_{list_fast_rec} :: 'a\ tree\ list \Rightarrow nat\\ &T_{list_fast_rec}\ ts\\ &= (\textbf{let}\ us = filter\ (\lambda t.\ t \neq \langle\rangle)\ ts\\ &\quad \textbf{in}\ |ts| +\\ &\qquad (\textbf{if}\ us = []\ \textbf{then}\ 0\\ &\qquad\ \textbf{else}\ 5 \cdot |us| + T_{list_fast_rec}\ (map\ left\ us\ @\ map\ right\ us)))\end{aligned}$$

The following inductive property is an abstraction of the core of the termination argument of *list_fast_rec* above.

$$(\forall\, t{\in}\mathit{set}\ \mathit{ts}.\ t \neq \langle\rangle) \longrightarrow$$
$$(\sum_{t\leftarrow \mathit{ts}}\ k \cdot |t|) = (\sum_{t\leftarrow \mathit{map\ left\ ts}\ @\ \mathit{map\ right\ ts}}\ k \cdot |t|) + k \cdot |\mathit{ts}| \quad (11.18)$$

The suggestive notation $\sum x{\leftarrow}\mathit{xs}.\ f\ x$ abbreviates $\mathit{sum_list}\ (\mathit{map}\ f\ \mathit{xs})$.

Now we can state and prove a linear upper bound of $T_{\mathit{list_fast_rec}}$:

Theorem 11.9. $T_{\mathit{list_fast_rec}}\ \mathit{ts} \leq (\sum_{t\leftarrow \mathit{ts}}\ 7 \cdot |t| + 1)$

Proof. The proof is by induction on the size of *ts*, again using the measure function $\lambda t.\ 2 \cdot |t| + 1$ which decreases with recursive calls as we proved above. If $\mathit{ts} = []$ the claim is trivial. Now assume $\mathit{ts} \neq []$ and let $\mathit{us} = \mathit{filter}\ (\lambda t.\ t \neq \langle\rangle)\ \mathit{ts}$ and $\mathit{children} = \mathit{map\ left\ us}\ @\ \mathit{map\ right\ us}$.

$$\begin{aligned}
&T_{\mathit{list_fast_rec}}\ \mathit{ts} = T_{\mathit{list_fast_rec}}\ \mathit{children} + 5 \cdot |\mathit{us}| + |\mathit{ts}| \\
&\leq (\sum_{t\leftarrow \mathit{children}}\ 7 \cdot |t| + 1) + 5 \cdot |\mathit{us}| + |\mathit{ts}| && \text{by IH} \\
&= (\sum_{t\leftarrow \mathit{children}}\ 7 \cdot |t|) + 7 \cdot |\mathit{us}| + |\mathit{ts}| \\
&= (\sum_{t\leftarrow \mathit{us}}\ 7 \cdot |t|) + |\mathit{ts}| && \text{by (11.18)} \\
&\leq (\sum_{t\leftarrow \mathit{ts}}\ 7 \cdot |t|) + |\mathit{ts}| = (\sum_{t\leftarrow \mathit{ts}}\ 7 \cdot |t| + 1) && \square
\end{aligned}$$

11.6 Exercises

Exercise 11.1. Instead of first showing that Braun trees are almost complete, give a direct proof of $\mathit{braun}\ t \longrightarrow h\ t = \lceil \lg |t|_1 \rceil$ by first showing $\mathit{braun}\ t \longrightarrow 2^{h\ t} \leq 2 \cdot |t| + 1$ by induction.

Exercise 11.2. Show that function *bal* in Section 4.3.1 produces Braun trees: $n \leq |\mathit{xs}| \wedge \mathit{bal}\ n\ \mathit{xs} = (t,\ \mathit{zs}) \longrightarrow \mathit{braun}\ t$. (Isabelle hint: *bal* needs to be qualified as *Balance.bal*.)

Exercise 11.3. One can view Braun trees as tries (see Chapter 12) by indexing them not with a *nat* but a *bool list* where each bit tells us whether to go left or right (as explained at the start of Section 11.2). Function *nat_of* specifies the intended correspondence:

$$\mathit{nat_of} :: \mathit{bool\ list} \Rightarrow \mathit{nat}$$
$$\mathit{nat_of}\ [] = 1$$
$$\mathit{nat_of}\ (b\ \#\ \mathit{bs}) = 2 \cdot \mathit{nat_of}\ \mathit{bs} + (\textbf{if}\ b\ \textbf{then}\ 1\ \textbf{else}\ 0)$$

Define the counterparts of *lookup*1 and *update*1

$$\mathit{lookup_trie} :: \mathit{'a\ tree} \Rightarrow \mathit{bool\ list} \Rightarrow \mathit{'a}$$
$$\mathit{update_trie} :: \mathit{bool\ list} \Rightarrow \mathit{'a} \Rightarrow \mathit{'a\ tree} \Rightarrow \mathit{'a\ tree}$$

and prove their correctness:

$$braun\ t \wedge nat_of\ bs \in \{1..|t|\} \longrightarrow$$
$$lookup_trie\ t\ bs = lookup1\ t\ (nat_of\ bs)$$

$$update_trie\ bs\ x\ t = update1\ (nat_of\ bs)\ x\ t$$

Exercise 11.4. Function *del_lo* is defined with the help of function *merge*. Define a recursive function $del_lo2 :: 'a\ tree \Rightarrow 'a\ tree$ without recourse to any auxiliary function and prove $del_lo2\ t = del_lo\ t$.

Exercise 11.5. Let *lh*, the "left height", compute the length of the left spine of a tree. Prove that the left height of a Braun tree is equal to its height: $braun\ t \longrightarrow lh\ t = h\ t$

Exercise 11.6. Show that the running time of *size_fast* is quadratic in the height of the tree: Define the running time functions T_{diff} and T_{size_fast} (taking 0 time in the base cases) and prove $T_{size_fast}\ t \leq (h\ t)^2$.

Exercise 11.7. Prove correctness of function *braun_of_naive* defined in Section 11.5.2: $list\ (braun_of_naive\ x\ n) = replicate\ n\ x$.

Bibliographic Remarks

Braun trees were investigated by Braun and Rem [67] and later, in a functional setting, by Hoogerwoord [31] who coined the term "Braun tree". Section 11.5 is partly based on work by Okasaki [61]. The whole chapter is based on work by Nipkow and Sewell [58].

12

Tries

A **trie** is a search tree where keys are strings, i.e. lists. A trie can be viewed as a tree-shaped finite automaton where the root is the start state. For example, the set of strings $\{\mathsf{a}, \mathsf{an}, \mathsf{can}, \mathsf{car}, \mathsf{cat}\}$ is encoded as the trie in Figure 12.1. The solid states are accepting, i.e. those nodes terminate the string leading to them.

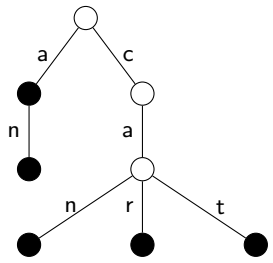


Fig. 12.1. A trie encoding $\{\mathsf{a}, \mathsf{an}, \mathsf{can}, \mathsf{car}, \mathsf{cat}\}$

12.1 Abstract Tries via Functions

A nicely abstract model of tries is the following type:

```
datatype 'a trie = Nd bool ('a ⇒ 'a trie option)
```

In a node $Nd\ b\ f$, b indicates if it is an accepting node and f maps characters to sub-tries.

There is no invariant, i.e. the invariant is simply *True*: there are no ordering, balance or other requirements.

Below we use a variant of the function update notation $f(a := b)$:

$$f(a \mapsto b) \equiv f(a := Some\ b)$$

This is how the ADT *Set* is implemented by means of tries:

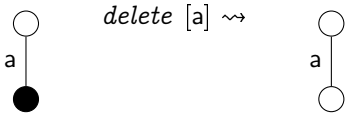
```
empty :: 'a trie
empty = Nd False (λ_. None)

isin :: 'a trie ⇒ 'a list ⇒ bool
isin (Nd b _) [] = b
isin (Nd _ m) (k # xs)
= (case m k of None ⇒ False | Some t ⇒ isin t xs)

insert :: 'a list ⇒ 'a trie ⇒ 'a trie
insert [] (Nd _ m) = Nd True m
insert (x # xs) (Nd b m)
= (let s = case m x of None ⇒ empty | Some t ⇒ t
   in Nd b (m(x ↦ insert xs s)))

delete :: 'a list ⇒ 'a trie ⇒ 'a trie
delete [] (Nd _ m) = Nd False m
delete (x # xs) (Nd b m)
= Nd b (case m x of None ⇒ m | Some t ⇒ m(x ↦ delete xs t))
```

The definitions are straightforward. But note that *delete* does not try to shrink the trie. For example:



Formally:

$$\begin{aligned}&delete\ [a]\ (Nd\ False\ [a \mapsto Nd\ True\ (\lambda_.\ None)])\\ &= Nd\ False\ [a \mapsto Nd\ False\ (\lambda_.\ None)]\end{aligned}$$

where $[x \mapsto t] \equiv (\lambda_.\ None)(x \mapsto t)$. The resulting trie is correct (it represents the empty set of strings) but could have been shrunk to $Nd\ False\ (\lambda_.\ None)$.

Functional Correctness

For the correctness proof we take a lazy approach and define the abstraction function in a trivial manner via *isin*:

$$set :: 'a\ trie \Rightarrow 'a\ list\ set$$
$$set\ t = \{xs \mid isin\ t\ xs\}$$

Correctness of *empty* and *isin* ($set\ empty = \{\}$ and $isin\ t\ xs = (xs \in set\ t)$) are trivial, correctness of insertion and deletion are easily proved by induction:

$$set\ (insert\ xs\ t) = set\ t \cup \{xs\}$$
$$set\ (delete\ xs\ t) = set\ t - \{xs\}$$

This simple model of tries leads to simple correctness proofs but is computationally inefficient because of the function space in $'a \Rightarrow 'a\ trie\ option$. In principle, any representation of this function space, for example by some search tree, works. However, such a generic theory is relatively complex. Hence we restrict ourselves to binary tries when exploring more efficient implementations of tries.

12.2 Binary Tries

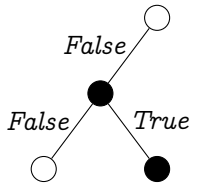
A **binary trie** is a trie over the alphabet *bool*. That is, binary tries represent sets of *bool list*. More concretely, every node has two children:

datatype $trie = Lf \mid Nd\ bool\ (trie \times trie)$

A binary trie, for example

$$Nd\ False\ (Nd\ True\ (Nd\ False\ (Lf, Lf), Nd\ True\ (Lf, Lf)), Lf)$$

can be visualized like this:



*Lf*s are not shown at all. The edge labels indicated that *False* refers to the left and *True* to the right child. This convention is encoded in the following auxiliary functions selecting from and modifying pairs:

```
sel2 :: bool ⇒ 'a × 'a ⇒ 'a
sel2 b (a1, a2) = (if b then a2 else a1)

mod2 :: ('a ⇒ 'a) ⇒ bool ⇒ 'a × 'a ⇒ 'a × 'a
mod2 f b (a1, a2) = (if b then (a1, f a2) else (f a1, a2))
```

The implementation of the *Set* interface is shown in Figure 12.2. Function *delete* shrinks a non-accepting *Nd* if both children have become empty.

```
empty :: trie
empty = Lf

isin :: trie ⇒ bool list ⇒ bool
isin Lf ks = False
isin (Nd b lr) ks = (case ks of [] ⇒ b | k # ks' ⇒ isin (sel2 k lr) ks')

insert :: bool list ⇒ trie ⇒ trie
insert [] Lf = Nd True (Lf, Lf)
insert [] (Nd _ lr) = Nd True lr
insert (k # ks) Lf = Nd False (mod2 (insert ks) k (Lf, Lf))
insert (k # ks) (Nd b lr) = Nd b (mod2 (insert ks) k lr)

delete :: bool list ⇒ trie ⇒ trie
delete _ Lf = Lf
delete ks (Nd b lr)
= (case ks of [] ⇒ node False lr
   | k # ks' ⇒ node b (mod2 (delete ks') k lr))

node b lr = (if ¬ b ∧ lr = (Lf, Lf) then Lf else Nd b lr)
```

Fig. 12.2. Implementation of *Set* via binary tries

Functional Correctness

For the correctness proof we take the same lazy approach as above:

```
set_trie :: trie ⇒ bool list set
set_trie t = {xs | isin t xs}
```

We are also lazy in that we set the invariant to *True*. A more precise invariant would express that the tries are minimal, i.e. cannot be shrunk. See Exercise 12.2. It turns out that the correctness properties do not require this more precise invariant.

The two non-trivial correctness properties are

$$set_trie\ (insert\ xs\ t) = set_trie\ t \cup \{xs\} \tag{12.1}$$

$$set_trie\ (delete\ xs\ t) = set_trie\ t - \{xs\} \tag{12.2}$$

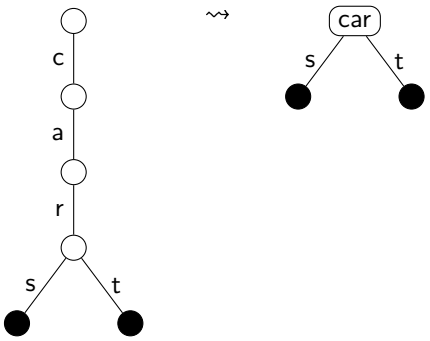
are simple consequences of the following inductive properties:

$$isin\ (insert\ xs\ t)\ ys = (xs = ys \vee isin\ t\ ys)$$

$$isin\ (delete\ xs\ t)\ ys = (xs \neq ys \wedge isin\ t\ ys)$$

12.3 Binary Patricia Tries

Tries can contain long branches without branching. These can be contracted by storing the branch directly in the start node. The result is called a **Patricia trie**. The following figure shows the contraction of a trie into a Patricia trie:



This is the data type of binary Patricia tries:

datatype $trieP = LfP \mid NdP\ (bool\ list)\ bool\ (trieP \times trieP)$

The implementation of the *Set* interface via binary Patricia tries is shown in Figure 12.3.

Functional Correctness

This is an exercise in stepwise data refinement. We have already proved that *trie* implements *Set* via an abstraction function. Now we map *trieP* back

```
emptyP = LfP

isinP LfP ks = False
isinP (NdP ps b lr) ks
= (let n = |ps|
   in if ps = take n ks
      then case drop n ks of [] ⇒ b | k # x ⇒ isinP (sel2 k lr) x
      else False)

insertP ks LfP = NdP ks True (LfP, LfP)
insertP ks (NdP ps b lr)
= (case split ks ps of (_, [], []) ⇒ NdP ps True lr
   | (qs, [], p # ps') ⇒
      let t = NdP ps' b lr
      in NdP qs True (if p then (LfP, t) else (t, LfP))
   | (_, k # ks', []) ⇒ NdP ps b (mod2 (insertP ks') k lr)
   | (qs, k # ks', _ # ps') ⇒
      let tp = NdP ps' b lr; tk = NdP ks' True (LfP, LfP)
      in NdP qs False (if k then (tp, tk) else (tk, tp)))

deleteP ks LfP = LfP
deleteP ks (NdP ps b lr) =
(case split ks ps of
  (qs, ks', p # ps') ⇒ NdP ps b lr
| (qs, k # ks', []) ⇒ nodeP ps b (mod2 (deleteP ks') k lr)
| (qs, [], []) ⇒ nodeP ps False lr)

nodeP ps b lr = (if ¬ b ∧ lr = (LfP, LfP) then LfP else NdP ps b lr)
```

Fig. 12.3. Implementation of *Set* via binary Patricia tries

to *trie* via another abstraction function. Afterwards the overall correctness follows trivially by composing the two abstraction functions.

The abstraction function *abs_trieP* is defined via an auxiliary function that prefixes a trie with a bit list:

```
prefix_trie :: bool list ⇒ trie ⇒ trie
prefix_trie [] t = t
prefix_trie (k # ks) t
= (let t' = prefix_trie ks t
   in Nd False (if k then (Lf, t') else (t', Lf)))
```

```
abs_trieP :: trieP ⇒ trie
abs_trieP LfP = Lf
abs_trieP (NdP ps b (l, r))
= prefix_trie ps (Nd b (abs_trieP l, abs_trieP r))
```

Again we take a lazy approach and set the invariant on *trieP* to *True*.

Correctness of *emptyP* is trivial. Correctness of the remaining operations is proved by induction and requires a number of supporting inductive lemmas which we display before the corresponding correctness property.

Correctness of *isinP*:

$$\begin{aligned}&isin\ (prefix_trie\ ps\ t)\ ks = (ps = take\ |ps|\ ks \wedge isin\ t\ (drop\ |ps|\ ks))\\&isinP\ t\ ks = isin\ (abs_trieP\ t)\ ks\end{aligned} \tag{12.3}$$

Correctness of *insertP*:

$$\begin{aligned}&prefix_trie\ ks\ (Nd\ True\ (Lf, Lf)) = insert\ ks\ Lf\\&insert\ ps\ (prefix_trie\ ps\ (Nd\ b\ lr)) = prefix_trie\ ps\ (Nd\ True\ lr)\\&insert\ (ks\ @\ ks')\ (prefix_trie\ ks\ t) = prefix_trie\ ks\ (insert\ ks'\ t)\\&prefix_trie\ (ps\ @\ qs)\ t = prefix_trie\ ps\ (prefix_trie\ qs\ t)\\&split\ ks\ ps = (qs, ks', ps') \longrightarrow\\&ks = qs\ @\ ks' \wedge ps = qs\ @\ ps' \wedge\\&(ks' \neq [] \wedge ps' \neq [] \longrightarrow hd\ ks' \neq hd\ ps')\\&abs_trieP\ (insertP\ ks\ t) = insert\ ks\ (abs_trieP\ t)\end{aligned} \tag{12.4}$$

Correctness of *deleteP*:

$$\begin{aligned}&delete\ xs\ (prefix_trie\ xs\ (Nd\ b\ (l, r)))\\&= (\textbf{if}\ (l, r) = (Lf, Lf)\ \textbf{then}\ Lf\ \textbf{else}\ prefix_trie\ xs\ (Nd\ False\ (l, r)))\\&delete\ (xs\ @\ ys)\ (prefix_trie\ xs\ t)\\&= (\textbf{if}\ delete\ ys\ t = Lf\ \textbf{then}\ Lf\ \textbf{else}\ prefix_trie\ xs\ (delete\ ys\ t))\\&abs_trieP\ (deleteP\ ks\ t) = delete\ ks\ (abs_trieP\ t)\end{aligned} \tag{12.5}$$

It is now trivial to obtain the correctness of the *trieP* implementation of sets. The abstraction function is simply the composition of the two abstraction abstraction functions: $set_trieP = set_trie \circ abs_trieP$. The required correctness properties (ignoring *emptyP* and *isinP*)

$$\begin{aligned}&set_trieP\ (insertP\ xs\ t) = set_trieP\ t \cup \{xs\}\\&set_trieP\ (deleteP\ xs\ t) = set_trieP\ t - \{xs\}\end{aligned}$$

are trivial compositions of (12.1)/(12.2) and (12.4)/(12.5).

12.4 Exercises

Exercise 12.1. Rework the above theory of binary tries as follows. Eliminate the *bool* argument from constructor *Nd* by replacing *Nd* by two constructors representing *Nd True* and *Nd False*.

Exercise 12.2. Define the invariant *invar* that characterizes fully shrunk binary tries, i.e. tries where every non-*Lf* trie represents a non-empty set. Note that a trie represents the empty set if it does not contain any node *Nd True* _. Prove that *insert* and *delete* maintain the invariant.

Bibliographic Remarks

Tries were first sketched by De La Briandais [15] and described in more detail by Fredkin [20] who coined their name based on the word reTRIEval. However, "trie" is usually pronounced like "try" rather than "tree" to avoid confusion. Patricia tries are due to Morrison [50].

13

Huffman's Algorithm

Huffman's algorithm [32] is a simple and elegant procedure for constructing a binary tree with minimum weighted path length—a measure of cost that considers both the lengths of the paths from the root to the leaf nodes and the weights associated with the leaf nodes. The algorithm's main application is data compression: By equating leaf nodes with characters, and weights with character frequencies, we can use it to derive optimum binary codes. A *binary code* is a map from characters to non-empty sequences of bits.

This chapter presents Huffman's algorithm and its optimality proof. In a slight departure from the rest of this book, the emphasis is more on graphical intuitions and less on rigorous logical arguments.

13.1 Binary Codes

Suppose we want to encode strings over a finite source alphabet as sequences of bits. Fixed-length codes like ASCII are simple and fast, but they generally waste space. If we know the frequency w_a of each source symbol a, we can save space by using shorter code words 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 code word for a.

As an example, consider the string '*abacabad*'. Encoding it with the code

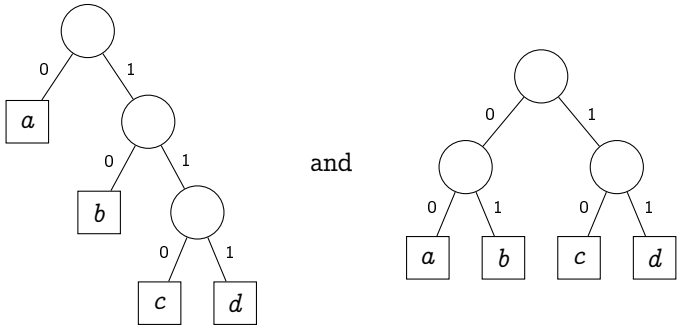
$$C_1 = \{a \mapsto 0,\ b \mapsto 10,\ c \mapsto 110,\ d \mapsto 111\}$$

gives the 14-bit code word 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*'. 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 the same string.

Binary codes can be represented by binary trees. For example, the trees



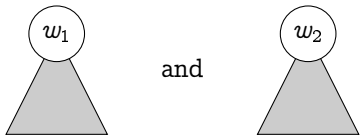
correspond to C_1 and C_2. The code word for a given symbol can be obtained as follows: Start at the root and descend toward the leaf node associated with the symbol one node at a time; emit a 0 whenever the left child of the current node is chosen and a 1 whenever the right child is chosen. The generated sequence of 0s and 1s is the code word.

To avoid ambiguities, we require that only leaf nodes are labeled with symbols. This ensures that no code word is a prefix of another. Moreover, it is sufficient to consider only full binary trees (trees whose inner nodes all have two children), because any node with only one child can advantageously be eliminated by removing it and letting the child take its parent's place.

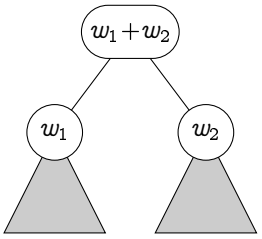
Each node in a code tree is assigned a *weight*. For a leaf node, the weight is the frequency of its symbol; for an inner node, it is the sum of the weights of its subtrees. In diagrams, we often annotate the nodes with their weights.

13.2 The Algorithm

Huffman's algorithm is a very simple procedure for constructing an optimum code tree for specified symbol frequencies: First, create a list of leaf nodes, one for each symbol in the alphabet, taking the given symbol frequencies as node weights. The nodes must be sorted in increasing order of weight. Second, pick the two trees



with the lowest weights and insert the tree



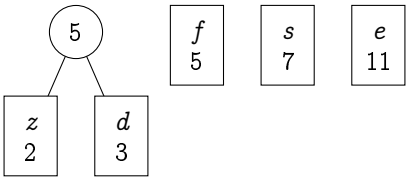
into the list so as to keep it ordered. Finally, repeat the process until only one tree is left in the list.

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:

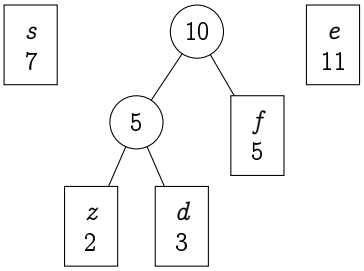
1.



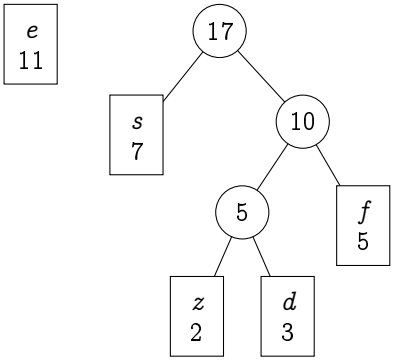
2.



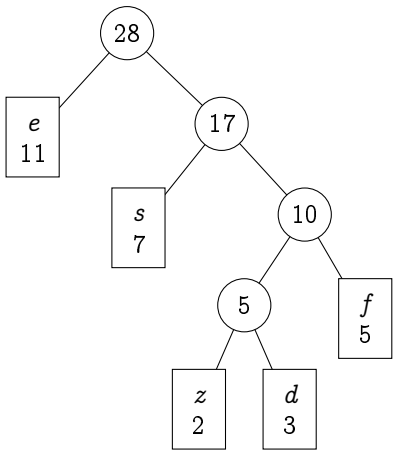
3.



4.



5.



The resulting tree is optimum for the given frequencies.

13.3 The Implementation

The functional implementation of the algorithm relies on the following type:

```
datatype 'a tree = Leaf nat 'a | Node nat ('a tree) ('a tree)
```

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 *Node w* t_1 t_2, where t_1 and t_2 are the left and right subtrees and *w* caches the sum of the weights of t_1 and t_2. The *cachedWeight* function extracts the weight stored in a node:

```
cachedWeight :: 'a tree ⇒ nat
cachedWeight (Leaf w _) = w
cachedWeight (Node w _ _) = w
```

The implementation builds on two additional auxiliary functions. The first one, *uniteTrees*, combines two trees by adding an inner node above them:

```
uniteTrees :: 'a tree ⇒ 'a tree ⇒ 'a tree
uniteTrees t1 t2 = Node (cachedWeight t1 + cachedWeight t2) t1 t2
```

The second function, *insortTree*, inserts a tree into a list sorted by cached weight, preserving the sort order:

```
insortTree :: 'a tree ⇒ 'a tree list ⇒ 'a tree list
insortTree u [] = [u]
insortTree u (t # ts)
= (if cachedWeight u ≤ cachedWeight t then u # t # ts
   else t # insortTree u ts)
```

The main function that implements Huffman's algorithm follows:

```
huffman :: 'a tree list ⇒ 'a tree
huffman [t] = t
huffman (t₁ # t₂ # ts) = huffman (insortTree (uniteTrees t₁ t₂) ts)
```

The function should initially be invoked with a non-empty list of leaf nodes sorted by weight. It repeatedly unites the first two trees of the list it receives as argument until a single tree is left.

13.4 Basic Auxiliary Functions Needed for the Proof

This section introduces basic concepts such as alphabet, consistency and optimality, which are needed to state the correctness and optimality of Huffman's algorithm. The next section introduces more specialized functions that arise in the proof.

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

```
alphabet :: 'a tree ⇒ 'a set
alphabet (Leaf _ a) = {a}
alphabet (Node _ t₁ t₂) = alphabet t₁ ∪ alphabet t₂
```

A tree is *consistent* if for each inner node the alphabets of the two subtrees are disjoint. Intuitively, this means that a symbol occurs in at most one leaf node. Consistency is a sufficient condition for δ_a (the length of the code word for a) to be uniquely defined. This well-formedness property appears as an assumption in many of the lemmas. The definition follows:

```
consistent :: 'a tree ⇒ bool
consistent (Leaf _ _) = True
consistent (Node _ t1 t2)
= (alphabet t1 ∩ alphabet t2 = {} ∧ consistent t1 ∧ consistent t2)
```

The *depth* of a symbol (which we wrote as δ_a above) is the length of the path from the root to that symbol, or equivalently the length of the code word for the symbol:

```
depth :: 'a tree ⇒ 'a ⇒ nat
depth (Leaf _ _) _ = 0
depth (Node _ t1 t2) a
= (if a ∈ alphabet t1 then depth t1 a + 1
   else if a ∈ alphabet t2 then depth t2 a + 1 else 0)
```

By convention, symbols that do not occur in the tree or that occur at the root of a one-node tree are given a depth of 0. If a symbol occurs in several leaf nodes (of an inconsistent tree), the depth is arbitrarily defined in terms of the leftmost node labeled with that symbol.

The *height* of a tree is the length of the longest path from the root to a leaf node, or equivalently the length of the longest code word:

```
height :: 'a tree ⇒ nat
height (Leaf _ _) = 0
height (Node _ t1 t2) = max (height t1) (height t2) + 1
```

The *frequency* of a symbol (which we wrote as w_a above) is the sum of the weights attached to the leaf nodes labeled with that symbol:

```
freq :: 'a tree ⇒ 'a ⇒ nat
freq (Leaf w a) b = (if b = a then w else 0)
freq (Node _ t1 t2) b = freq t1 b + freq t2 b
```

For consistent trees, the sum comprises at most one nonzero term. The frequency is then the weight of the leaf node labeled with the symbol, or 0 if there is no such node.

Two trees are *comparable* if they have the same alphabet and symbol frequencies. This is an important concept, because it allows us to state not

only that the tree constructed by Huffman's algorithm is optimal but also that it has the expected alphabet and frequencies.

The *weight* function returns the weight of a tree:

```
weight :: 'a tree ⇒ nat
weight (Leaf w _) = w
weight (Node _ t1 t2) = weight t1 + weight t2
```

In the *Node* case, we ignore the weight cached in the node and instead compute the tree's weight recursively.

The *cost* (or *weighted path length*) of a consistent tree is the sum

$$\sum_{a \in \mathit{alphabet}\ t} \mathit{freq}\ t\ a \cdot \mathit{depth}\ t\ a$$

which we wrote as $\sum_a w_a \delta_a$ above. It is defined recursively by

```
cost :: 'a tree ⇒ nat
cost (Leaf _ _) = 0
cost (Node _ t1 t2) = weight t1 + cost t1 + weight t2 + cost t2
```

A tree is optimum if and only if its cost is not greater than that of any comparable tree:

```
optimum :: 'a tree ⇒ bool
optimum t
= (∀ u. consistent u ⟶
        alphabet t = alphabet u ⟶
        freq t = freq u ⟶ cost t ≤ cost u)
```

Tree functions are readily generalized to lists of trees, or *forests*. For example, the alphabet of a forest is defined as the union of the alphabets of its trees. The forest generalizations have a subscript '$_F$' attached to their name (e.g. $\mathit{alphabet}_F$).

13.5 Other Functions Needed for the Proof

The optimality proof needs to interchange nodes in trees, to replace a two-leaf subtree with weights w_1 and w_2 by a single leaf of weight $w_1 + w_2$ and

vice versa, and to refer to the two symbols with the lowest frequencies. These concepts are represented by four functions: *swapSyms*, *swapFourSyms*, *mergeSibling*, *splitLeaf* and *minima*.

The interchange function *swapSyms* takes a three t and two symbols a, b, and exchanges them:

```
swapSyms :: 'a tree ⇒ 'a ⇒ 'a ⇒ 'a tree
swapSyms t a b = swapLeaves t (freq t a) a (freq t b) b
```

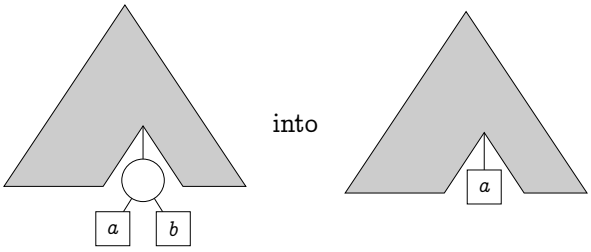
The following lemma captures the intuition that to minimize the cost, more frequent symbols should be encoded using fewer bits than less frequent ones:

Lemma 13.1. *consistent* $t \wedge a \in$ *alphabet* $t \wedge b \in$ *alphabet* $t \wedge$ *freq* $t\ a \leq$ *freq* $t\ b \wedge$
depth $t\ a \leq$ *depth* $t\ b \longrightarrow$
cost (*swapSyms* $t\ a\ b$) $\leq$ *cost* t

The four-way symbol interchange function *swapFourSyms* 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's and d's positions. A naive definition of this function would be *swapSyms* (*swapSyms* $t\ a\ c$) $b\ d$. This naive definition fails in the face of aliasing: If $a = d$, but $b \neq c$, then *swapFourSyms* $a\ b\ c\ d$ would leave a in b's position. Instead, we use this definition:

```
swapFourSyms :: 'a tree ⇒ 'a ⇒ 'a ⇒ 'a ⇒ 'a ⇒ 'a tree
swapFourSyms t a b c d
= (if a = d then swapSyms t b c
    else if b = c then swapSyms t a d
         else swapSyms (swapSyms t a c) b d)
```

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



The frequency of a in the resulting tree is the sum of the original frequencies of a and b. The function is defined by the equations

```
mergeSibling :: 'a tree ⇒ 'a ⇒ 'a tree

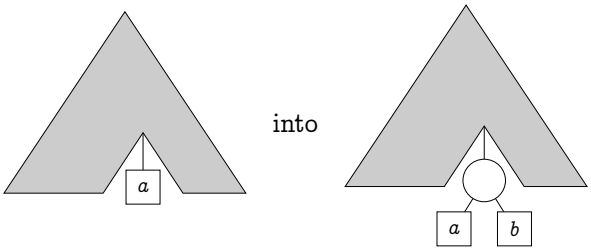
mergeSibling (Leaf w_b b) _ = Leaf w_b b
mergeSibling (Node w (Leaf w_b b) (Leaf w_c c)) a
= (if a = b ∨ a = c then Leaf (w_b + w_c) a
   else Node w (Leaf w_b b) (Leaf w_c c))
mergeSibling (Node w (Node v va vb) t_2) a
= Node w (mergeSibling (Node v va vb) a) (mergeSibling t_2 a)
mergeSibling (Node w t_1 (Node v va vb)) a
= Node w (mergeSibling t_1 a) (mergeSibling (Node v va vb) a)
```

The *sibling* function returns the label of the node that is the (left or right) sibling of the node labeled with the given symbol a in tree t. If a is not in t's alphabet or it occurs in a node with no sibling leaf, we simply return a. This gives us the nice property that if t is consistent, then $sibling\ t\ a \neq a$ if and only if a has a sibling. The definition, which is omitted here, distinguishes the same cases as *mergeSibling*.

Using the *sibling* function, we can state that merging two sibling leaves with weights w_a and w_b decreases the cost by $w_a + w_b$:

Lemma 13.2. $consistent\ t \wedge sibling\ t\ a \neq a \longrightarrow$
$cost\ (mergeSibling\ t\ a) + freq\ t\ a + freq\ t\ (sibling\ t\ a) = cost\ t$

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 *splitLeaf* with w_a and w_b such that $freq\ t\ a = w_a + w_b$. The definition follows:

```
splitLeaf :: 'a tree ⇒ nat ⇒ 'a ⇒ nat ⇒ 'a ⇒ 'a tree

splitLeaf (Leaf w_c c) w_a a w_b b
= (if c = a then Node w_c (Leaf w_a a) (Leaf w_b b) else Leaf w_c c)
splitLeaf (Node w t_1 t_2) w_a a w_b b
```

$= \mathit{Node}\ w\ (\mathit{splitLeaf}\ t_1\ w_a\ a\ w_b\ b)\ (\mathit{splitLeaf}\ t_2\ w_a\ a\ w_b\ b)$

Splitting a leaf with weight $w_a + w_b$ into two sibling leaves with weights w_a and w_b increases the cost by $w_a + w_b$:

Lemma 13.3. $\mathit{consistent}\ t \wedge a \in \mathit{alphabet}\ t \wedge \mathit{freq}\ t\ a = w_a + w_b \longrightarrow \mathit{cost}\ (\mathit{splitLeaf}\ t\ w_a\ a\ w_b\ b) = \mathit{cost}\ t + w_a + w_b$

Finally, the *minima* predicate expresses that two symbols a, b have the lowest frequencies in the tree t and that $\mathit{freq}\ t\ a \leq \mathit{freq}\ t\ b$:

$$
\begin{aligned}
&\mathit{minima} :: 'a\ \mathit{tree} \Rightarrow 'a \Rightarrow 'a \Rightarrow \mathit{bool}\\
&\mathit{minima}\ t\ a\ b\\
&= (a \in \mathit{alphabet}\ t \wedge b \in \mathit{alphabet}\ t \wedge a \neq b \wedge\\
&\quad (\forall c{\in}\mathit{alphabet}\ t.\\
&\qquad c \neq a \longrightarrow c \neq b \longrightarrow \mathit{freq}\ t\ a \leq \mathit{freq}\ t\ c \wedge \mathit{freq}\ t\ b \leq \mathit{freq}\ t\ c))
\end{aligned}
$$

13.6 The Key Lemmas and Theorems

It is easy to prove that the tree returned by Huffman's algorithm preserves the alphabet, consistency and symbol frequencies of the original forest:

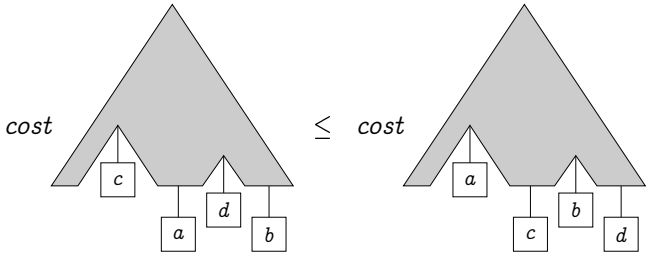
$\mathit{ts} \neq [] \longrightarrow \mathit{alphabet}\ (\mathit{huffman}\ \mathit{ts}) = \mathit{alphabet}_F\ \mathit{ts}$

$\mathit{consistent}_F\ \mathit{ts} \wedge \mathit{ts} \neq [] \longrightarrow \mathit{consistent}\ (\mathit{huffman}\ \mathit{ts})$

$\mathit{ts} \neq [] \longrightarrow \mathit{freq}\ (\mathit{huffman}\ \mathit{ts})\ a = \mathit{freq}_F\ \mathit{ts}\ a$

The main difficulty is to prove the optimality of the tree constructed by Huffman's algorithm. We need three lemmas before we can present the optimality theorem.

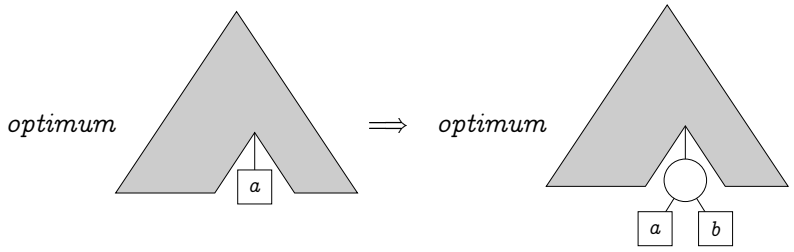
First, if a and b are minima, and c and d are at the very bottom of the tree, then exchanging a and b with c and d does not increase the tree's cost. Graphically, we have



Lemma 13.4. *consistent t* $\wedge$ *minima t a b* $\wedge$
c $\in$ *alphabet t* $\wedge$ *d* $\in$ *alphabet t* $\wedge$
depth t c $=$ *height t* $\wedge$ *depth t d* $=$ *height t* $\wedge$ *c* $\neq$ *d* $\longrightarrow$
cost (*swapFourSyms t a b c d*) $\leq$ *cost t*

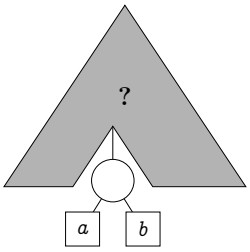
Proof. The proof is by case analysis on $a = c$, $a = d$, $b = c$, and $b = d$. The cases are easy to prove by expanding the definition of *swapFourSyms* and applying Lemma 13.1. □

The tree *splitLeaf t* w_a *a* w_b *b* is optimum if *t* is optimum, under a few assumptions, notably that *freq t a* $= w_a + w_b$. Graphically:



Lemma 13.5. *consistent t* $\wedge$ *optimum t* $\wedge$
a $\in$ *alphabet t* $\wedge$ *b* $\notin$ *alphabet t* $\wedge$ *freq t a* $= w_a + w_b$ $\wedge$
($\forall$ *c*$\in$ *alphabet t*. $w_a \leq$ *freq t c* $\wedge$ $w_b \leq$ *freq t c*) $\longrightarrow$
optimum (*splitLeaf t* w_a *a* w_b *b*)

Proof. We assume that *t*'s cost is less than or equal to that of any other comparable tree *v* and show that *splitLeaf t* w_a *a* w_b *b* has a cost less than or equal to that of any other comparable tree *u*. For the nontrivial case where *height t* > 0, it is easy to prove that there must be two symbols *c* and *d* occurring in sibling nodes at the very bottom of *u*. From *u* we construct the tree *swapFourSyms u a b c d* in which the minima *a* and *b* are siblings:



The question mark reminds us that we know nothing specific about *u*'s structure. Merging *a* and *b* gives a tree comparable with *t*, which we can use to instantiate *v*:

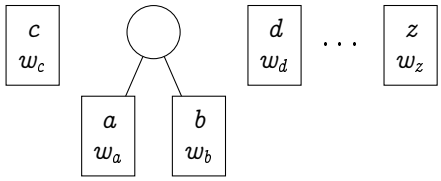
$$
\begin{aligned}
&\mathit{cost}\ (\mathit{splitLeaf}\ t\ a\ w_a\ b\ w_b) = \mathit{cost}\ t + w_a + w_b && \text{by Lemma 13.3}\\
&\leq \mathit{cost}\ (\mathit{mergeSibling}\ (\mathit{swapFourSyms}\ u\ a\ b\ c\ d)\ a) + w_a + w_b\\
& && \text{by optimality assumption}\\
&= \mathit{cost}\ (\mathit{swapFourSyms}\ u\ a\ b\ c\ d) && \text{by Lemma 13.2}\\
&\leq \mathit{cost}\ u && \text{by Lemma 13.4}\ \square
\end{aligned}
$$

Once it has combined two lowest-weight trees using *uniteTrees*, Huffman's algorithm does not 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.

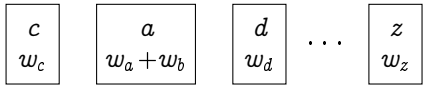
Lemma 13.6. $\mathit{consistent}_F\ ts \wedge ts \neq [] \wedge a \in \mathit{alphabet}_F\ ts \wedge$
$\mathit{freq}_F\ ts\ a = w_a + w_b \longrightarrow$
$\mathit{splitLeaf}\ (\mathit{huffman}\ ts)\ w_a\ a\ w_b\ b$
$= \mathit{huffman}\ (\mathit{splitLeaf}_F\ ts\ w_a\ a\ w_b\ b)$

The proof is by straightforward induction on the length of the forest *ts*.

As a consequence of this commutativity lemma, applying Huffman's algorithm on a forest of the form



gives the same result as applying the algorithm on the "flat" forest

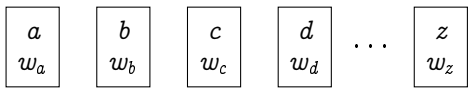


followed by splitting the leaf node a into two nodes a and b with frequencies w_a, w_b. The lemma effectively provides a way to flatten the forest at each step of the algorithm.

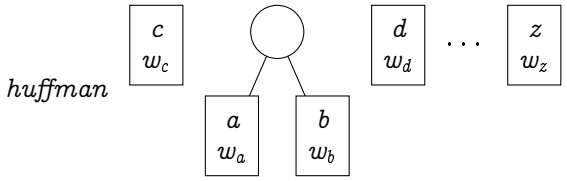
This leads us to our main result.

Theorem 13.7. $\mathit{consistent}_F\ ts \wedge \mathit{height}_F\ ts = 0 \wedge \mathit{sortedByWeight}\ ts \wedge$
$ts \neq [] \longrightarrow$
$\mathit{optimum}\ (\mathit{huffman}\ ts)$

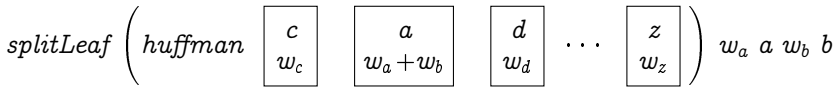
Proof. The proof is by induction on the length of *ts*. The assumptions ensure that *ts* is of the form



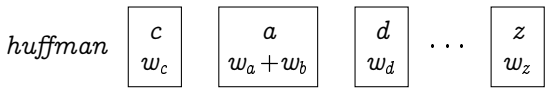
with $w_a \leq w_b \leq w_c \leq w_d \leq \cdots \leq w_z$. If *ts* consists of a single node, the node has cost 0 and is therefore optimum. If *ts* has length 2 or more, the first step of the algorithm leaves us with a term such as



In the diagram, we put the newly created tree at position 2 in the forest; in general, it could be anywhere. By Lemma 13.6, the above tree equals



To prove that this tree is optimum, it suffices by Lemma 13.5 to show that



is optimum, which follows from the induction hypothesis. □

In summary, we have established that the *huffman* program, which constitutes a functional implementation of Huffman's algorithm, constructs a binary tree that represents an optimal binary code for the specified alphabet and frequencies.

13.7 Chapter Notes

The sorted list of trees constitutes a simple priority queue (Part III). The time complexity of Huffman's algorithm is quadratic in the size n of this queue. By using a binary search to implement *insortTree*, we can obtain an $O(n \lg n)$ imperative implementation. An $O(n)$ implementation is possible by maintaining two queues, one containing the unprocessed leaf nodes and the other containing the combined trees [43].

Huffman's algorithm was invented by David Huffman [32]. The proof above was inspired by Knuth's informal argument [43]. This chapter's text is based on a published article [8]. An alternative formal proof, developed using Coq, is due to Théry [73].

Knuth [42] presented an alternative, more abstract view of Huffman's algorithm as a "Huffman algebra." Could his approach help simplify our proof? The most tedious steps above concerned splitting nodes, merging siblings, and

swapping symbols. These steps would still be necessary, as the algebraic approach seems restricted to abstracting over the arithmetic reasoning, which is not very difficult in the first place. On the other hand, with Knuth's approach, perhaps the proof would gain in elegance.

Part III

Priority Queues

14

Priority Queues

A **priority queue** of linearly ordered elements is like a multiset where one can insert arbitrary elements and remove minimal elements. Its specification as an ADT is show in Figure 14.1 where $Min_mset\ m \equiv Min\ (set_mset\ m)$ and Min yields the minimal element of a finite set of linearly ordered elements.

ADT $Priority_Queue =$

interface
$empty :: 'q$
$insert :: 'a \Rightarrow 'q \Rightarrow 'q$
$del_min :: 'q \Rightarrow 'q$
$get_min :: 'q \Rightarrow 'a$

abstraction $mset :: 'q \Rightarrow 'a\ multiset$
invariant $invar :: 'q \Rightarrow bool$

specification

$invar\ empty$	$(empty\text{-}inv)$
$mset\ empty = \{\!\{\}\!\}$	$(empty)$
$invar\ q \longrightarrow invar\ (insert\ x\ q)$	$(insert\text{-}inv)$
$invar\ q \longrightarrow mset\ (insert\ x\ q) = mset\ q + \{\!\{x\}\!\}$	$(insert)$
$invar\ q \wedge mset\ q \neq \{\!\{\}\!\} \longrightarrow invar\ (del_min\ q)$	$(del_min\text{-}inv)$
$invar\ q \wedge mset\ q \neq \{\!\{\}\!\} \longrightarrow$ $mset\ (del_min\ q) = mset\ q - \{\!\{get_min\ q\}\!\}$	(del_min)
$invar\ q \wedge mset\ q \neq \{\!\{\}\!\} \longrightarrow get_min\ q = Min_mset\ (mset\ q)$	(get_min)

Fig. 14.1. ADT *Priority_Queue*

Mergeable priority queues (see Figure 14.2) provide an additional function *merge* (sometimes: *meld* or *union*) with the obvious functionality.

ADT *Priority_Queue_Merge* = *Priority_Queue* +

interface
merge :: $'q \Rightarrow 'q \Rightarrow 'q$

specification
$invar\ q_1 \wedge invar\ q_2 \longrightarrow invar\ (merge\ q_1\ q_2)$
$invar\ q_1 \wedge invar\ q_2 \longrightarrow mset\ (merge\ q_1\ q_2) = mset\ q_1 + mset\ q_2$

Fig. 14.2. ADT *Priority_Queue_Merge*

Our priority queues are simplified. The more general version contains elements that are pairs of some item and its priority. In that case a priority queue can be viewed as a set, but because we have dropped the item we need to view a priority queue as a multiset. In imperative implementations, priority queues frequently also provide an operation *decrease_key*: given some direct reference to an element in the priority queue, decrease the element's priority. This is not completely straightforward in a functional language. Lammich and Nipkow [46] present an implementation, a Priority Search Tree.

Exercise 14.1. Give a list-based implementation of mergeable priority queues with constant-time *get_min* and *del_min*. Verify the correctness of your implementation w.r.t. *Priority_Queue_Merge*.

14.1 Heaps

A popular implementation technique for priority queues are **heaps**, i.e. trees where the minimal element in each subtree is at the root:

```
heap :: 'a tree ⇒ bool

heap ⟨⟩ = True
heap ⟨l, m, r⟩
= ((∀x∈set_tree l ∪ set_tree r. m ≤ x) ∧ heap l ∧ heap r)
```

Function *mset_tree* extracts the multiset of elements from a tree:

```
mset_tree :: 'a tree ⇒ 'a multiset

mset_tree ⟨⟩ = {#}
mset_tree ⟨l, a, r⟩ = {#a#} + mset_tree l + mset_tree r
```

When verifying a heap-based implementation of priority queues the invariant *invar* and the abstraction function *mset* in the ADT *Priority_Queue* are instantiated by *heap* and *mset_tree*. The correctness proofs need to talk about both multisets and (because of the *heap* invariant) sets of elements in a heap. We will only show the relevant multiset properties because the set properties follow easily via $set_mset\ (mset_tree\ t) = set_tree\ t$.

Both *empty* and *get_min* have obvious implementations:

$$empty = \langle\rangle$$

$$get_min = value$$

where $value\ \langle _,\ a,\ _\rangle = a$. If a heap-based implementation provides a *merge* function (e.g. skew heaps in Chapter 22), then *insert* and *del_min* can be defined like this:

$$insert\ x\ t = merge\ \langle\langle\rangle,\ x,\ \langle\rangle\rangle\ t$$

$$\begin{aligned} &del_min\ \langle\rangle = \langle\rangle \\ &del_min\ \langle l,\ _,\ r\rangle = merge\ l\ r \end{aligned}$$

Note that the following tempting definition of *merge* is functionally correct but leads to very unbalanced heaps:

$$\begin{aligned} &merge\ \langle\rangle\ t = t \\ &merge\ t\ \langle\rangle = t \\ &merge\ (\langle l_1,\ a_1,\ r_1\rangle =: t_1)\ (\langle l_2,\ a_2,\ r_2\rangle =: t_2) \\ &= (\textbf{if}\ a_1 \le a_2\ \textbf{then}\ \langle l_1,\ a_1,\ merge\ r_1\ t_2\rangle\ \textbf{else}\ \langle l_2,\ a_2,\ merge\ t_1\ r_2\rangle) \end{aligned}$$

Many of the more advanced implementations of heaps focus on improving this merge function. We will see examples of this in the next chapter on leftist heaps, as well as in the chapters on skew heaps and pairing heaps.

Exercise 14.2. Show functional correctness of the above definition of *merge* (w.r.t. *Priority_Queue_Merge*) and prove functional correctness of the implementations of *insert* and *del_min* (w.r.t. *Priority_Queue*).

Bibliographic Remarks

The idea of the heap goes back to Williams [77] who also coined the name.

15

Leftist Heaps

Leftist heaps are heaps in the sense of Section 14.1 and implement mergeable priority queues. The key idea is to maintain the invariant that at each node the minimal height of the right child is $\leq$ that of the left child. We represent leftist heaps as augmented trees that store the minimal height in every node:

```
type_synonym 'a lheap = ('a × nat) tree

mht :: 'a lheap ⇒ nat
mht ⟨⟩ = 0
mht ⟨_, (_, n), _⟩ = n
```

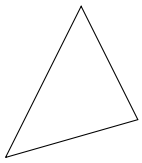
There are two invariants: the standard *heap* invariant (on augmented trees)

```
heap :: ('a × 'b) tree ⇒ bool
heap ⟨⟩ = True
heap ⟨l, (m, _), r⟩
= ((∀x∈set_tree l ∪ set_tree r. m ≤ x) ∧ heap l ∧ heap r)
```

and the structural **leftist tree** invariant that requires that the minimal height of the right child is no bigger than that of the left child (and that the minimal height information in the node is correct):

```
ltree :: 'a lheap ⇒ bool
ltree ⟨⟩ = True
ltree ⟨l, (_, n), r⟩
= (mh r ≤ mh l ∧ n = mh r + 1 ∧ ltree l ∧ ltree r)
```

Thus a tree is a leftist tree if for every subtree the right spine is a shortest path from the root to a leaf. Pictorially:



Now remember $2^{mh\ t} \leq |t|_1$, i.e. $mh\ t \leq \lg\ |t|_1$. Because the expensive operations on leftist heaps descend along the right spine, this means that their running time is logarithmic in the size of the heap.

Exercise 15.1. An alternative definition of leftist tree is via the length of the right spine of the tree:

$rank :: 'a\ tree \Rightarrow nat$
$rank\ \langle\rangle = 0$
$rank\ \langle _,\ _,\ r\rangle = rank\ r + 1$

Prove that $ltree\ t \longrightarrow rank\ t = mh\ t$.

15.1 Implementation of ADT *Priority_Queue_Merge*

The key operation is *merge*:

```
merge :: 'a lheap ⇒ 'a lheap ⇒ 'a lheap
merge ⟨⟩ t = t
merge t ⟨⟩ = t
merge (⟨l₁, (a₁, n₁), r₁⟩ =: t₁) (⟨l₂, (a₂, n₂), r₂⟩ =: t₂)
= (if a₁ ≤ a₂ then node l₁ a₁ (merge r₁ t₂)
   else node l₂ a₂ (merge t₁ r₂))

node :: 'a lheap ⇒ 'a ⇒ 'a lheap ⇒ 'a lheap
node l a r
= (let mhl = mht l; mhr = mht r
   in if mhr ≤ mhl then ⟨l, (a, mhr + 1), r⟩
      else ⟨r, (a, mhl + 1), l⟩)
```

Termination of *merge* can be proved either by the sum of the sizes of the two arguments (which goes down with every call) or by the lexicographic product of the two size measures: either the first argument becomes smaller or it stays unchanged and the second argument becomes smaller.

As shown in Section 14.1, once we have *merge*, the other operations are easily definable. We repeat their definitions simply because this chapter employs augmented rather than ordinary trees:

```
empty :: 'a lheap
empty = ⟨⟩

get_min :: 'a lheap ⇒ 'a
get_min ⟨_, (a, _), _⟩ = a

insert :: 'a ⇒ 'a lheap ⇒ 'a lheap
insert x t = merge ⟨⟨⟩, (x, 1), ⟨⟩⟩ t

del_min :: 'a lheap ⇒ 'a lheap
del_min ⟨⟩ = ⟨⟩
del_min ⟨l, _, r⟩ = merge l r
```

15.2 Correctness

The above implementation is proved correct w.r.t. the ADT *Priority_Queue_Merge* where

```
mset_tree :: ('a × 'b) tree ⇒ 'a multiset
mset_tree ⟨⟩ = {#}
mset_tree ⟨l, (a, _), r⟩ = {#a#} + mset_tree l + mset_tree r

invar t = (heap t ∧ ltree t)
```

Correctness of *get_min* follows directly from the heap invariant:

$$heap\ t \wedge t \neq \langle\rangle \longrightarrow get_min\ t = Min\ (set_tree\ t)$$

From the following inductive lemmas about *merge*

$$mset_tree\ (merge\ t_1\ t_2) = mset_tree\ t_1 + mset_tree\ t_2$$

$$ltree\ l \wedge ltree\ r \longrightarrow ltree\ (merge\ l\ r)$$

$$heap\ l \wedge heap\ r \longrightarrow heap\ (merge\ l\ r)$$

correctness of *insert* and *del_min* follow easily:

$$mset_tree\ (insert\ x\ t) = mset_tree\ t + \{\!\{x\}\!\}$$
$$mset_tree\ (del_min\ t) = mset_tree\ t - \{\!\{get_min\ t\}\!\}$$
$$ltree\ t \longrightarrow ltree\ (insert\ x\ t)$$
$$heap\ t \longrightarrow heap\ (insert\ x\ t)$$
$$ltree\ t \longrightarrow ltree\ (del_min\ t)$$
$$heap\ t \longrightarrow heap\ (del_min\ t)$$

Of course the above proof (ignoring the *ltree* part) works for any mergeable priority queue implemented as a heap.

15.3 Running Time Analysis

The running time functions are shown in Appendix B.4. By induction on the computation of merge we obtain

$$ltree\ l \wedge ltree\ r \longrightarrow T_{merge}\ l\ r \leq mh\ l + mh\ r + 1$$

With $2^{mh\ t} \leq |t|_1$ it follows that

$$ltree\ l \wedge ltree\ r \longrightarrow T_{merge}\ l\ r \leq \lg |l|_1 + \lg |r|_1 + 1 \tag{15.1}$$

which implies logarithmic bounds for insertion and deletion:

$$ltree\ t \longrightarrow T_{insert}\ x\ t \leq \lg |t|_1 + 3$$
$$ltree\ t \longrightarrow T_{del_min}\ t \leq 2 \cdot \lg |t|_1 + 1$$

The derivation of the bound for insertion is trivial but the one for deletion uses a little lemma, assuming *ltree* t:

$$\begin{aligned}
&T_{del_min}\ t = T_{merge}\ l\ r + 1 \\
&\leq \lg |l|_1 + \lg |r|_1 + 2 && \text{using (15.1)} \\
&\leq 2 \cdot \lg |t|_1 + 1 && \text{because } \lg x + \lg y + 1 < 2 \cdot \lg (x + y) \\
& && \text{if } 0 < x \text{ and } 0 < y
\end{aligned}$$

Bibliographic Remarks

Leftist heaps were invented by Crane [14]. Another version of leftist trees, based on weight rather than height, was introduced by Cho and Sahni [12, 68].

16

Priority Queues via Braun Trees

In Chapter 11 we introduced Braun trees and showed how to implement arrays. In the current chapter we show how to implement priority queues by means of Braun trees. Because Braun trees have logarithmic height this guarantees logarithmic running times for insertion and deletion. Remember that every node $\langle l, x, r\rangle$ in a Braun tree satisfies $|l| \in \{|r|, |r| + 1\}$ $(*)$.

16.1 Implementation of ADT *Priority_Queue*

We follow the heap approach in Section 14.1. Functions *empty*, *get_min*, *heap* and *mset_tree* are defined as in that section.

Insertion and deletion maintain the Braun tree property $(*)$ by inserting into the right (and possibly smaller) child, deleting from the left (and possibly larger) child, and swapping children to reestablish $(*)$.

Insertion is straightforward and clearly maintains both the heap and the Braun tree property:

```
insert :: 'a ⇒ 'a tree ⇒ 'a tree
insert a ⟨⟩ = ⟨⟨⟩, a, ⟨⟩⟩
insert a ⟨l, x, r⟩
= (if a < x then ⟨insert x r, a, l⟩ else ⟨insert a r, x, l⟩)
```

To delete the minimal (i.e. root) element from a tree, extract the leftmost element from the tree and let it sift down to its correct position in the tree in the manner of heapsort:

```
del_min :: 'a tree ⇒ 'a tree
del_min ⟨⟩ = ⟨⟩
del_min ⟨⟨⟩, x, r⟩ = ⟨⟩
del_min ⟨l, x, r⟩ = (let (y, l') = del_left l in sift_down r y l')

del_left :: 'a tree ⇒ 'a × 'a tree
del_left ⟨⟨⟩, x, r⟩ = (x, r)
del_left ⟨l, x, r⟩ = (let (y, l') = del_left l in (y, ⟨r, x, l'⟩))

sift_down :: 'a tree ⇒ 'a ⇒ 'a tree ⇒ 'a tree
sift_down ⟨⟩ a _ = ⟨⟨⟩, a, ⟨⟩⟩
sift_down ⟨⟨⟩, x, _⟩ a ⟨⟩
= (if a ≤ x then ⟨⟨⟩, x, ⟨⟩⟩, a, ⟨⟩⟩ else ⟨⟨⟨⟩, a, ⟨⟩⟩, x, ⟨⟩⟩)
sift_down (⟨l1, x1, r1⟩ =: t1) a (⟨l2, x2, r2⟩ =: t2)
= (if a ≤ x1 ∧ a ≤ x2 then ⟨t1, a, t2⟩
   else if x1 ≤ x2 then ⟨sift_down l1 a r1, x1, t2⟩
        else ⟨t1, x2, sift_down l2 a r2⟩)
```

In the first two equations for *sift_down*, the Braun tree property guarantees that the "_" arguments must be empty trees if the pattern matches.

16.2 Correctness

We outline the correctness proofs for *insert* and *del_min* by presenting the key lemmas. Correctness of *insert* is straightforward:

$|insert\ x\ t| = |t| + 1$

$mset_tree\ (insert\ x\ t) = \{\!\{x\}\!\} + mset_tree\ t$

$braun\ t \longrightarrow braun\ (insert\ x\ t)$

$heap\ t \longrightarrow heap\ (insert\ x\ t)$

Correctness of *del_min* builds on analogous correctness lemmas for the auxiliary functions:

$del_left\ t = (x, t') \wedge t \neq \langle\rangle \longrightarrow mset_tree\ t = \{\!\{x\}\!\} + mset_tree\ t'$

$del_left\ t = (x, t') \wedge t \neq \langle\rangle \wedge heap\ t \longrightarrow heap\ t'$

$del_left\ t = (x, t') \wedge t \neq \langle\rangle \longrightarrow |t| = |t'| + 1$

$del_left\ t = (x, t') \wedge t \neq \langle\rangle \wedge braun\ t \longrightarrow braun\ t'$

$$braun\ \langle l,\ a,\ r\rangle \longrightarrow |sift_down\ l\ a\ r| = |l| + |r| + 1$$

$$braun\ \langle l,\ a,\ r\rangle \longrightarrow braun\ (sift_down\ l\ a\ r)$$

$$braun\ \langle l,\ a,\ r\rangle \longrightarrow$$
$$mset_tree\ (sift_down\ l\ a\ r) = \{\!\{a\}\!\} + (mset_tree\ l + mset_tree\ r)$$

$$braun\ \langle l,\ a,\ r\rangle \wedge heap\ l \wedge heap\ r \longrightarrow heap\ (sift_down\ l\ a\ r)$$

$$braun\ t \longrightarrow braun\ (del_min\ t)$$

$$heap\ t \wedge braun\ t \longrightarrow heap\ (del_min\ t)$$

$$braun\ t \longrightarrow |del_min\ t| = |t| - 1$$

$$braun\ t \wedge t \neq \langle\rangle \longrightarrow$$
$$mset_tree\ (del_min\ t) = mset_tree\ t - \{\!\{get_min\ t\}\!\}$$

Bibliographic Remarks

Our implementation of priority queues via Braun trees is due to Paulson [64] who credits it to Okasaki.

17

Binomial Heaps

Binomial heaps are another common implementation of mergeable priority queues, which supports efficient ($O(\log n)$) *insert*, *get_min*, *del_min*, and *merge* operations.

The basic building blocks of a binomial heap are *binomial trees*, which are defined recursively as follows. A binomial tree of rank r is a node with r children of ranks $r-1, \ldots, 0$, in that order. Figure 17.1 shows an example binomial tree.

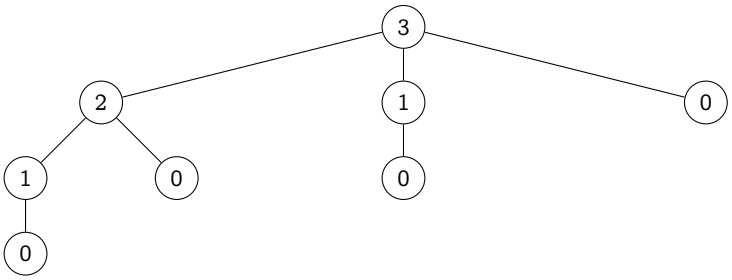


Fig. 17.1. A binomial tree of rank 3. The node labels depict the rank of each node. A node of rank r has child nodes of ranks $r-1, \ldots, 0$.

To define binomial trees in Isabelle, we first define a more general datatype:

```
datatype 'a tree = Node nat 'a ('a tree list)
```

Apart from the list of children, a node stores a rank and an element. As usual, we use a shortcut notation for nodes:

```
⟨r, x, ts⟩ ≡ Node r x ts
```

This datatype contains all binomial trees, but also some non-binomial trees. To carve out the binomial trees, we define an invariant, which reflects the informal definition above:

```
invar_btree :: 'a tree ⇒ bool
invar_btree ⟨r, _, ts⟩
= ((∀ t∈set ts. invar_btree t) ∧ map rank ts = rev [0..<r])
```

Additionally, we require the heap property, i.e., that the root element of each subtree is a minimal element in that subtree:

```
invar_otree :: 'a tree ⇒ bool
invar_otree ⟨_, x, ts⟩ = (∀ t∈set ts. invar_otree t ∧ x ≤ root t)
```

Thus, a *binomial tree* is a tree that satisfies both the structural and the heap invariant. The two invariants are combined in a single predicate:

```
invar_tree :: 'a tree ⇒ bool
invar_tree t = (invar_btree t ∧ invar_otree t)
```

A *binomial heap* is a list of binomial trees

```
type_synonym 'a heap = 'a tree list
```

with strictly ascending rank:

```
invar :: 'a heap ⇒ bool
invar ts
= ((∀ t∈set ts. invar_tree t) ∧ sorted_wrt (<) (map rank ts))
```

Note that *sorted_wrt* states that a list is sorted wrt. the specified relation, here $(<)$. It is defined in Appendix A.

17.1 Size

The following functions return the multiset of elements in a binomial tree and in a binomial heap:

```
mset_tree :: 'a tree ⇒ 'a multiset
mset_tree ⟨_, a, ts⟩ = {#a#} + ∑# (image_mset mset_tree (mset ts))
mset_heap :: 'a heap ⇒ 'a multiset
mset_heap ts = ∑# (image_mset mset_tree (mset ts))
```

Most operations on binomial heaps are linear in the length of the heap. To show that the length is bounded by the number of heap elements, we first observe that the number of elements in a binomial tree is already determined by its rank. A binomial tree of rank r has 2^r nodes:

$$invar_btree\ t \longrightarrow |mset_tree\ t| = 2^{rank\ t}$$

This property is proved by induction on the tree structure. A tree of rank 0 has one element, and a tree of rank $r+1$ has subtrees of rank $0, 1, \ldots, r$. By the induction hypothesis, these have $2^0, 2^1, \ldots, 2^r$ elements, i.e., $2^{r+1} - 1$ elements together. Including the element at the root, there are 2^{r+1} elements.

The length of a binomial heap is bounded logarithmically in the number of its elements:

$$invar\ ts \longrightarrow |ts| \leq \lg\ (|mset_heap\ ts| + 1)$$

To prove this, recall that the heap ts is strictly sorted by rank. Thus, we can underestimate the ranks of the trees in ts by $0, 1, \ldots, |ts| - 1$. This means that they must have at least $2^0, 2^1, \ldots, 2^{|ts|-1}$ elements, i.e., at least $2^{|ts|} - 1$ elements together, which yields the desired bound.

17.2 Implementation of ADT *Priority_Queue*

17.2.1 Empty and Emptiness Test

Obviously, the empty list is a binomial heap with no elements, and a binomial heap is empty only if it is the empty list:

```
invar []
mset_heap [] = {#}
(mset_heap ts = {#}) = (ts = [])
```

17.2.2 Insertion

A crucial property of binomial trees is that we can link two binomial trees of rank r to form a binomial tree of rank $r + 1$, simply by prepending one tree as the first child of the other. To preserve the heap property, we add the tree with the bigger root element below the tree with the smaller root element. This *linking* of trees is illustrated in Figure 17.2. Formally:

$$\begin{aligned}&\mathit{link} :: \mathit{'a\ tree} \Rightarrow \mathit{'a\ tree} \Rightarrow \mathit{'a\ tree}\\&\mathit{link}\ (\langle r,\ x_1,\ ts_1\rangle =: t_1)\ (\langle r',\ x_2,\ ts_2\rangle =: t_2)\\&= (\textbf{if}\ x_1 \leq x_2\ \textbf{then}\ \langle r + 1,\ x_1,\ t_2\ \#\ ts_1\rangle\ \textbf{else}\ \langle r + 1,\ x_2,\ t_1\ \#\ ts_2\rangle)\end{aligned}$$

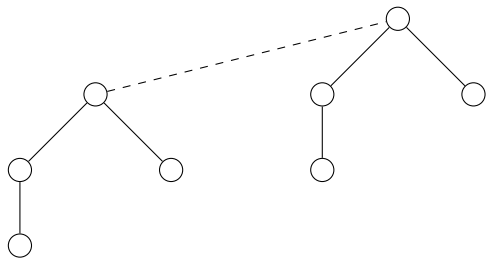


Fig. 17.2. Linking two binomial trees of rank 2 to form a binomial tree of rank 3, by linking the left tree as first child of the right tree, as indicated by the dashed line. We assume that the root element of the left tree is greater than or equal to the root element of the right tree, such that the heap property is preserved.

By case distinction, we can easily prove that *link* preserves the invariant and that the resulting tree contains the elements of both arguments.

$$\begin{aligned}&\mathit{invar_tree}\ t_1 \wedge \mathit{invar_tree}\ t_2 \wedge \mathit{rank}\ t_1 = \mathit{rank}\ t_2 \longrightarrow\\&\mathit{invar_tree}\ (\mathit{link}\ t_1\ t_2)\end{aligned}$$

$$\mathit{mset_tree}\ (\mathit{link}\ t_1\ t_2) = \mathit{mset_tree}\ t_1 + \mathit{mset_tree}\ t_2$$

The link operation forms the basis of inserting a tree into a heap: if the heap does not contain a tree with the same rank, we can simply insert the tree at the correct position in the heap. Otherwise, we merge the two trees and recursively insert the result. For our purposes, we can additionally assume that the rank of the tree to be inserted is smaller than or equal to the lowest rank in the heap, which saves us a case in the following definition:

```
ins_tree :: 'a tree ⇒ 'a heap ⇒ 'a heap
ins_tree t [] = [t]
ins_tree t₁ (t₂ # ts)
= (if rank t₁ < rank t₂ then t₁ # t₂ # ts
    else ins_tree (link t₁ t₂) ts)
```

Invariant preservation and functional correctness of *ins_tree* is easily proved by induction using the respective properties for *link*:

```
invar_tree t ∧ invar ts ∧ (∀ t'∈set ts. rank t ≤ rank t') −→
invar (ins_tree t ts)
mset_heap (ins_tree t ts) = mset_tree t + mset_heap ts
```

A single element is inserted as a one-element (rank 0) tree:

```
insert :: 'a ⇒ 'a heap ⇒ 'a heap
insert x ts = ins_tree ⟨0, x, []⟩ ts
```

The above definition meets the specification for insert required by the *Priority_Queue* ADT:

```
invar t −→ invar (insert x t)
mset_heap (insert x t) = {|x|} + mset_heap t
```

17.2.3 Merging

Recall the merge algorithm used in top-down merge sort (Section 2.4). It merges two sorted lists by repeatedly taking the smaller list head. We use a similar idea for merging two heaps: if the rank of one list's head is strictly smaller, we choose it. If both ranks are equal, we link the two heads and insert the resulting tree into the merged remaining heaps. Thus, the resulting heap will be strictly ordered by rank. Formally:

```
merge :: 'a heap ⇒ 'a heap ⇒ 'a heap
merge ts₁ [] = ts₁
merge [] (v # va) = v # va
merge (t₁ # ts₁ =: h₁) (t₂ # ts₂ =: h₂)
= (if rank t₁ < rank t₂ then t₁ # merge ts₁ h₂
    else if rank t₂ < rank t₁ then t₂ # merge h₁ ts₂
         else ins_tree (link t₁ t₂) (merge ts₁ ts₂))
```

The *merge* function can be regarded as an algorithm for adding two sparse binary numbers. This intuition is explored in Exercise 17.2.

We show that the merge operation preserves the invariant and unites the elements:

$$invar\ ts_1 \wedge invar\ ts_2 \longrightarrow invar\ (merge\ ts_1\ ts_2)$$
$$mset_heap\ (merge\ ts_1\ ts_2) = mset_heap\ ts_1 + mset_heap\ ts_2$$

The proof is straightforward, except for preservation of the binomial heap invariant. We first show that merging two heaps does not decrease the lowest rank in these heaps. This ensures that prepending the head with smaller rank to the merged remaining heaps results in a sorted heap. Moreover, when we link two heads of equal rank, this ensures that the linked tree's rank is smaller than or equal to the ranks in the merged remaining trees, as required by the *ins_tree* function. We phrase this property as preservation of lower rank bounds, i.e., a lower rank bound of both heaps is still a lower bound for the merged heap:

$$t' \in set\ (merge\ ts_1\ ts_2) \wedge (\forall\, t_1 \in set\ ts_1.\ rank\ t < rank\ t_1) \wedge$$
$$(\forall\, t_2 \in set\ ts_2.\ rank\ t < rank\ t_2) \longrightarrow$$
$$rank\ t < rank\ t'$$

The proof is by straightforward induction, relying on an analogous bounding lemma for *ins_tree*.

17.2.4 Finding a Minimal Element

For a binomial tree, the root node always contains a minimal element. Unfortunately, there is no such property for the whole heap—the minimal element may be at the root of any of the heap's trees. To get a minimal element from a non-empty heap, we look at all root nodes:

```
get_min :: 'a heap ⇒ 'a
get_min [t] = root t
get_min (t # v # va) = min (root t) (get_min (v # va))
```

Correctness of this operation is proved by a simple induction:

$$mset_heap\ ts \neq \{\!\}\ \wedge\ invar\ ts \longrightarrow$$
$$get_min\ ts = Min_mset\ (mset_heap\ ts)$$

17.2.5 Deleting a Minimal Element

To delete a minimal element, we first need to find one and then remove it. Removing the root node of a tree with rank r leaves us with a list of its

children, which are binomial trees of ranks $r-1,\ldots,0$. Reversing this list yields a valid binomial heap, which we merge with the remaining trees in the original heap:

```
del_min :: 'a heap ⇒ 'a heap
del_min ts
= (case get_min_rest ts of (⟨_, _, ts₁⟩, ts₂) ⇒ merge (rev ts₁) ts₂)
```

Here, the auxiliary function *get_min_rest* splits a heap into a tree with minimal root element, and the remaining trees.

```
get_min_rest :: 'a heap ⇒ 'a tree × 'a heap
get_min_rest [t] = (t, [])
get_min_rest (t # v # va)
= (let (t', ts') = get_min_rest (v # va)
    in if root t ≤ root t' then (t, v # va) else (t', t # ts'))
```

We prove that, for a non-empty heap, *del_min* preserves the invariant and deletes the minimal element:

$$ts \neq [] \wedge invar\ ts \longrightarrow invar\ (del_min\ ts)$$
$$ts \neq [] \longrightarrow mset_heap\ ts = mset_heap\ (del_min\ ts) + \{\!\{get_min\ ts\}\!\}$$

The proof is straightforward. For invariant preservation, the key is to show that *get_min_rest* returns a binomial tree and a binomial heap:

$$get_min_rest\ ts = (t', ts') \wedge ts \neq [] \wedge invar\ ts \longrightarrow invar_tree\ t'$$
$$get_min_rest\ ts = (t', ts') \wedge ts \neq [] \wedge invar\ ts \longrightarrow invar\ ts'$$

To show that we actually remove a minimal element, we show that *get_min_rest* selects the same tree as *get_min*:

$$ts \neq [] \wedge get_min_rest\ ts = (t', ts') \longrightarrow root\ t' = get_min\ ts$$

17.3 Running Time Analysis

The running time functions are shown in Appendix B.5. Intuitively, the operations are linear in the length of the heap, which in turn is logarithmic in the number of elements (cf. Section 17.1).

The running time analysis for *insert* is straightforward. The running time is dominated by *ins_tree*. In the worst case, it iterates over the whole heap, taking constant time per iteration. By straightforward induction, we show

$$T_{ins_tree}\ t\ ts \leq |ts| + 1$$

and thus

$$invar\ ts \longrightarrow T_{insert}\ x\ ts \leq \lg\ (|mset_heap\ ts| + 1) + 2$$

The running time analysis for merge is more interesting. In each recursion, we need constant time to compare the ranks. However, if the ranks are equal, we link the trees and insert them into the merger of the remaining heaps. In the worst case, this costs linear time in the length of the merger. A naive analysis would estimate $|merge\ ts_1\ ts_2| \leq |ts_1| + |ts_2|$, and thus yield a quadratic running time in the length of the heap.

However, we can do better: we observe that every link operation in *ins_tree* reduces the number of trees in the heap. Thus, over the whole merge, we can only have linearly many link operations in the combined size of both heaps.

To formalize this idea, we estimate the running time of *ins_tree* and *merge* together with the length of the result:

$$T_{ins_tree}\ t\ ts + |ins_tree\ t\ ts| = 2 + |ts|$$
$$|merge\ ts_1\ ts_2| + T_{merge}\ ts_1\ ts_2 \leq 2 \cdot (|ts_1| + |ts_2|) + 1$$

Both estimates can be proved by straightforward induction, and from the second estimate we easily derive a bound for *merge*:

$$invar\ ts_1 \wedge invar\ ts_2 \longrightarrow$$
$$T_{merge}\ ts_1\ ts_2 \leq 4 \cdot \lg\ (|mset_heap\ ts_1| + |mset_heap\ ts_2| + 1) + 1$$

From the bound for *merge*, we can easily derive a bound for *del_min*:

$$invar\ ts \wedge ts \neq [] \longrightarrow T_{del_min}\ ts \leq 6 \cdot \lg\ (|mset_heap\ ts| + 1) + 3$$

The only notable point is that we use a linear time bound for reversing a list, as explained in Section 1.5.1:

$$T_{rev} :: 'a\ list \Rightarrow nat$$
$$T_{rev}\ xs = |xs| + 1$$

17.4 Exercises

Exercise 17.1. A node in a tree is on level n if it needs n edges from the root to reach it. Define a function $nol::nat \Rightarrow 'a\ tree \Rightarrow nat$ such that $nol\ n\ t$ is the number of nodes on level n in tree t and show that a binomial tree of rank r has $\binom{r}{l}$ nodes on level l:

$$invar_btree\ t \longrightarrow nol\ l\ t = rank\ t\ choose\ l$$

Hint: You might want to prove separately that

$$\sum_{i=0}^{i<r} \binom{i}{n} = \binom{r}{n+1}$$

Exercise 17.2. Sparse binary numbers represent a binary number by a list of the positions of set bits, sorted in ascending order. Thus, the list [1, 3, 4] represents the number 11010. In general, $[p_1, \ldots, p_n]$ represents $2^{p_1} + \cdots + 2^{p_n}$.

Implement sparse binary numbers in Isabelle, using the type *nat list*.

1. Define a function *invar* :: *nat list* ⇒ *bool* that checks for strictly ascending bit positions, a function *num_of* :: *nat list* ⇒ *nat* that converts a sparse binary number to a natural number, and a function *add* :: *nat list* ⇒ *nat list* ⇒ *nat list* to add sparse binary numbers.
2. Show that your *add* function preserves the invariant and actually performs addition as far as *num_of* is concerned.
3. Define a running time function for *add* and show that it is linear in the list size.

Hint: The bit positions in sparse binary numbers are analogous to binomial trees of a certain rank in a binomial heap. The *add* function is implemented similar to the *merge* function, using a *carry* function to insert a bit position into a number (similar to *ins_tree*). Correctness and running time can be proved similarly.

Bibliographic Remarks

Binomial queues were invented by Vuillemin [75]. Functional implementations were given by King [40] and Okasaki [62]. A functional implementation was verified by Meis *et al.* [49]; a Java implementation, by Müller [51].

Part IV

Advanced Design and Analysis Techniques

18

Dynamic Programming

You probably have seen this function before:

$$
\begin{aligned}
&\mathit{fib} :: \mathit{nat} \Rightarrow \mathit{nat} \\
&\mathit{fib}\ 0 = 0 \\
&\mathit{fib}\ 1 = 1 \\
&\mathit{fib}\ (n + 2) = \mathit{fib}\ (n + 1) + \mathit{fib}\ n
\end{aligned}
$$

It computes the well-known Fibonacci numbers. You may also have noticed that calculating $\mathit{fib}\ 50$ already causes quite some stress for your computer and there is no hope for $\mathit{fib}\ 500$ to ever return a result.

This is quite unfortunate considering that there is a very simple imperative program to compute these numbers efficiently:

```
int fib(n) {
  int a = 0;
  int b = 1;
  for (i in 1..n) {
    int temp = b;
    b = a + b;
    a = temp;
  }
  return a;
}
```

So we seem to be caught in an adverse situation here: either we use a clear and elegeant definition of *fib* or we get an efficient but convoluted implementation for *fib*. Admittedly, we could just prove that both formulations are the

same function, and use whichever one is more suited for the task at hand. For *fib*, of course, it is trivial to define a functional analogue of the imperative program and to prove its correctness. However, doing this for all recursive functions we would like to define is tedious. Instead, this chapter will sketch a recipe that allows to define such recursive functions in the natural way, while still getting an efficient implementation "for free".

In the following, the Fibonacci function will serve as a simple example on which we can illustrate the idea. Next, we will show how to prove the correctness of the efficient implementation in an efficient way. Subsequently, we will discuss further details of the how approach and how it can be applied beyond *fib*. The chapter closes with the study of two famous (and archetypical) dynamic programming algorithms: the Bellman-Ford algorithm for finding shortest paths in weighted graphs and an algorithm due to Knuth for computing optimal binary search trees.

18.1 Memoization

Let us consider the tree of recursive calls that are issued when computing *fib* 5 (Fig. 18.1). We can see that the subtree for *fib* 3 is computed twice, and

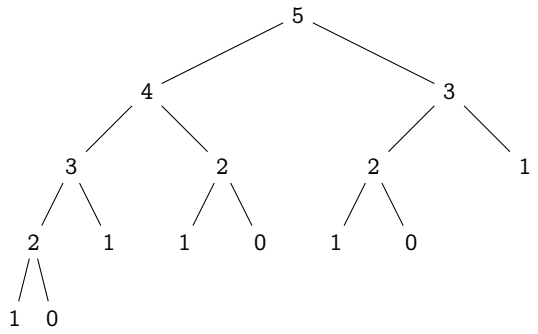


Fig. 18.1. Tree of the recursive call structure for *fib* 5

that the subtree for *fib* 2 is even computed three times. How can we avoid these repeated computations? A common solution is *memoization*: we store previous computation results in some kind of memory and consult it to potentially recall a memoized result before issuing another recursive computation. Here is a simple memoizing version of *fib* that implements the memory as a mapping of type $nat \rightharpoonup nat$ (the notation $m(k \mapsto v)$ is a shorthand for $m(k := Some\ v)$, see also Section 12.1):

```
fib₁ :: nat ⇒ (nat ⇀ nat) ⇒ nat × (nat ⇀ nat)
fib₁ 0 m = (0, m(0 ↦ 0))
fib₁ 1 m = (1, m(1 ↦ 1))
fib₁ (n + 2) m
= (let (i, m) = case m n of None ⇒ fib₁ n m | Some i ⇒ (i, m);
       (j, m) =
         case m (n + 1) of None ⇒ fib₁ (n + 1) m
         | Some j ⇒ (j, m)
   in (i + j, m(n + 2 ↦ i + j)))
```

And indeed, we can ask Isabelle to compute (via the *value* command) $fib_1\ 50$ or even $fib_1\ 500$ and we get the result within a split second.

However, we are not yet happy with this code. Carrying the memory around means a lot of additional weight for the definition of fib_1, and proving that this function computes the same value as *fib* is not completely trivial (how would you approach this?). Let us streamline the definition first by pulling out the reading and writing of memory into a function *memo* (for a type $'k$ of keys and a type $'v$ of values):

```
memo ::
  'k ⇒ (('k ⇀ 'v) ⇒ 'v × ('k ⇀ 'v))
       ⇒ ('k ⇀ 'v) ⇒ 'v × ('k ⇀ 'v)
memo k f m
= (case m k of None ⇒ let (v, m) = f m in (v, m(k ↦ v))
   | Some v ⇒ (v, m))

fib₂ :: nat ⇒ (nat ⇀ nat) ⇒ nat × (nat ⇀ nat)
fib₂ 0 = memo 0 (λm. (0, m))
fib₂ 1 = memo 1 (λm. (1, m))
fib₂ (n + 2)
= memo (n + 2)
   (λm. let (i, m) = fib₂ n m;
            (j, m) = fib₂ (n + 1) m
        in (i + j, m))
fib₂ (n + 2)
= memo (n + 2)
   (λm. let (i, m) = fib₂ n m; (j, y) = fib₂ (n + 1) m in (i + j, y))
```

This already looks a lot more like the original definition but it still has one problem: we have to thread the memory through the program explicitly. This can be become rather tedious for more complicated programs and diverges from the original shape of the program, complicating the proofs.

18.1.1 Enter the Monad

Let us examine the type of fib_2 more closely. We can read it as the type of a function that, given a natural number, returns a *computation*. Given an initial memory, it computes a pair of a result and an updated memory. We can capture this notion of "stateful" computations in a data type:

```
datatype ('s, 'a) state = State ('s ⇒ 'a × 's)
```

A value of type $('s, 'a)\ state$ represents a stateful computation that returns a result of type $'a$ and operates on states of type $'s$. The constant *run_state* forces the evaluation of a computation starting from some initial state:

```
run_state :: ('s, 'a) state ⇒ 's ⇒ 'a × 's
run_state (State f) s = f s
```

The advantage of this definition may not seem immediate. Its value only starts to show when we see how it allows us to *chain* stateful computations. To do so, we only need to define two constants: *return* to pack up a result in a computation, and *bind* to chain two computations after each other.

```
return :: 'a ⇒ ('s, 'a) state
return x = State (λs. (x, s))

bind :: ('s, 'a) state ⇒ ('a ⇒ ('s, 'b) state) ⇒ ('s, 'b) state
bind a f = State (λs. let (x, s) = run_state a s in run_state (f x) s)
```

We add a little syntax on top and write $\langle\!\langle x \rangle\!\rangle$ for *return* x, and $a \gg\!= f$ instead of *bind* a f. The "identity" computation $\langle\!\langle x \rangle\!\rangle$ simply leaves the given state unchanged and produces x as a result. The chained computation $a \gg\!= f$ starts with some state s, runs a on it to produce a pair of a result x and a new state s', and then evaluates $f\ x$ to produce another computation that is run on s'.

We have now seen how to pass state around but we are not yet able to interact with it. For this purpose we define *get* and *set* to retrieve and update the current state, respectively:

```
get :: ('s, 's) state
get = State (λs. (s, s))

set :: 's ⇒ ('s, unit) state
set s' = State (λ_. ((), s'))
```

Let us reformulate fib_2 with the help of these concepts:

```
memo₁ ::
  'k ⇒ ('k ⇀ 'v, 'v) state ⇒ ('k ⇀ 'v, 'v) state
memo₁ k a
= get ⤜
   (λm. case m k of
          None ⇒ a ⤜ (λv. set (m(k ↦ v)) ⤜ (λ_. ⟪v⟫))
        | Some x ⇒ ⟪x⟫)

fib₃ :: nat ⇒ (nat ⇀ nat, nat) state
fib₃ 0 = ⟪0⟫
fib₃ 1 = ⟪1⟫
fib₃ (n + 2)
= memo₁ (n + 2) (fib₃ n ⤜ (λi. fib₃ (n + 1) ⤜ (λj. ⟪i + j⟫))))
```

Can you see how we have managed to hide the whole handling of state behind the scenes? The only explicit interaction with the state is now happening inside of $memo_1$. This is sensible as this is the only place where we really want to recall a memoized result or to write a new value to memory.

While this is great, we still want to polish the definition further: the syntactic structure of the last case of fib_3 still does not match *fib* exactly. To this end, we lift function application $f\ x$ to the state monad:

```
(.) :: ('s, 'a ⇒ ('s, 'b) state) state ⇒ ('s, 'a) state ⇒ ('s, 'b) state
f_m . x_m = (f_m ⤜ (λf. x_m ⤜ (λx. f x)))
```

We can now spell out our final memoizing version of *fib* where $(.)$ replaces ordinary function applications in the original definition:

$$
\begin{aligned}
&\mathit{fib}_4 :: \mathit{nat} \Rightarrow (\mathit{nat} \rightharpoonup \mathit{nat}, \mathit{nat})\ \mathit{state}\\
&\mathit{fib}_4\ 0 = \langle\!\langle 0 \rangle\!\rangle\\
&\mathit{fib}_4\ 1 = \langle\!\langle 1 \rangle\!\rangle\\
&\mathit{fib}_4\ (n+2)\\
&= \mathit{memo}_1\ (n+2)\ (\langle\!\langle \lambda i.\ \langle\!\langle \lambda j.\ \langle\!\langle i + j \rangle\!\rangle \rangle\!\rangle \rangle\!\rangle\ .\ (\mathit{fib}_4\ n)\ .\ (\mathit{fib}_4\ (n+1)))
\end{aligned}
$$

You may wonder why we added that many additional computations in this last step. On the one hand, we have gained the advantage that we can now closely follow the syntactic structure of *fib* to prove that fib_4 is correct (notwithstanding that memo_1 will need a special treatment, of course). On the other hand, we can remove most of these additional computations in a final post-processing step.

18.1.2 Memoization and Dynamic Programming

Let us recap what we have seen so far in this chapter. We noticed that the naive recursive formulation of the Fibonacci numbers leads to a highly inefficient implementation. We then showed how to work around this problem by using memoization to obtain a structurally similar but efficient implementation. After all this, you may wonder why this chapter is titled *Dynamic Programming* and not *Memoization*.

Dynamic programming relies on two main principles. First, to find an optimal solution for a problem by computing it from optimal solutions for "smaller" instances of the same problem, i.e. *recursion*. Second, to *memoize* these solutions for smaller problems in, e.g. a table. Thus we could be bold and state:

$$\text{dynamic programming} = \text{recursion} + \text{memoization}$$

A common objection to this equation would be that memoization should be distinguished from *tabulation*. In this view, the former only computes "necessary" solutions for smaller sub-problems, while the latter just "blindly" builds solutions for sub-problems of increasing size, many of which might be unnecessary. The benefit of tabulation could be increased performance, for instance due to improved caching performance. We believe that this distinction is largely irrelevant to our approach. First, in this book we focus on asymptotically efficient solutions, not constant-factor optimizations. Second, in many dynamic programming algorithms memoization would actually compute solutions for the same set of sub-problems as tabulation does. No matter which of the two approaches is used in the implementation, the hard part is

usually to come up with a recursive solution that can efficiently make use of sub-problems in the first place.

There are problems, however, where clever tabulation instead of naive memoization is necessary to achieve an *asymptotically optimal solution* in terms of *memory consumption*. One instance of this is the Bellman-Ford algorithm presented in Section 18.4. On this example, we will show that our approach is also akin to tabulation. It can easily be introduced as a final "post-processing" step.

Some readers may have noticed that our optimized implementations of *fib* are not really optimal as they use a map for memoization. Indeed it is possible to swap in other memory implementations as long as they provide a *lookup* and an *update* method. One can even make use of imperative data structures like arrays. As this is not the focus of this book, the interested reader is referred to the literature that is provided at the end of this chapter. Here, we will just assume that the maps used for memoization are implemented as red-black trees (and Isabelle's code generator can be instructed to do so).

For the remainder of this chapter, we will first outline how to prove that fib_4 is correct. Then, we will sketch how to apply our approach of memoization beyond *fib*. Afterwards, we will study some prototypical examples of dynamic programming problems and show how to apply the above formula to them.

18.2 Correctness of Memoization

We now want to prove that fib_4 is correct. But what is it exactly that we want to prove? We surely want fib_4 to produce the same result as *fib* when run with an empty memory (where $empty \equiv \lambda x.\ None$):

$$fst\ (run_state\ (fib_4\ n)\ empty) = fib\ n \tag{18.1}$$

If we were to make a naive attempt at this prove, we would probably start with an induction on the computation of *fib* just to realize that the induction hypotheses are not strong enough to prove the recursion case, since they demand an empty memory. We can attempt generalization as a remedy:

$$fst\ (run_state\ (fib_4\ n)\ m) = fib\ n$$

However, this statement does not hold anymore for every memory m.

What do we need to demand from m? It should only memoize values that are *consistent* with *fib*:

```
type_synonym 'a mem = (nat ⇀ nat, 'a) state

cmem :: (nat ⇀ nat) ⇒ bool
cmem m = (∀ n∈dom m. m n = Some (fib n))

dom :: ('k ⇀ 'v) ⇒ 'k set
dom m = {a | m a ≠ None}
```

Note that, from now, we use the type *'a mem* to denote *memoized* values of type *'a* that have been "wrapped up" in our memoizing state monad. Using *cmem*, we can formulate a general notion of equivalence between a value v and its memoized version a, written $v \triangleright a$: starting from a consistent memory m, a should produce another consistent memory m', and the result v.

```
(▷) :: 'a ⇒ 'a mem ⇒ bool
v ▷ a
= (∀ m. cmem m ⟶
        (let (v', m') = run_state a m in v = v' ∧ cmem m'))
```

Thus we want to prove

$$fib\ n \triangleright fib_4\ n \tag{18.2}$$

via computation induction on n as above. For the base cases we need to prove statements of the form $v \triangleright \langle v \rangle$, a property which follows trivially after unfolding the involved definitions. For the induction case, we can unfold $fib_4\ (n + 2)$, and get rid of $memo_1$ by applying the following rule (which we instantiate with $a = fib_4\ n$):

$$fib\ n \triangleright a \longrightarrow fib\ n \triangleright memo_1\ n\ a \tag{18.3}$$

For the remainder of the proof, we now want to unfold $fib\ (n + 2)$ and then follow the syntactic structure of fib_4 and fib in lockstep. To do so, we need to find a proof rule for function application. That is, what do we need in order to prove $f\ x \triangleright f_m\ .\ x_m$? For starters, $x \triangleright x_m$ seems reasonable to demand. But what about f and f_m? If f has type $'a \Rightarrow 'b$, then f_m is of type $('a \Rightarrow 'b\ mem)\ mem$. Intuitively, we want to state something along these lines:

> "f_m is a *memoized* function that, when applied to a value x, yields a memoized value that is equivalent to $f\ x$".

This goes beyond what we can currently express with $(\triangleright)$ as $v \triangleright a$ merely states that "a is a memoized value equivalent to v". What we need is more liberty in our choice of equivalence. That is, we want to use statements $v \triangleright_R a$, with the meaning: "a is a memoized value that is related to v by R". The formal definition is analogous to $(\triangleright)$ (and $(\triangleright) = (\triangleright_{(=)})$):

$$
\begin{aligned}
&(\cdot \triangleright \cdot) :: 'a \Rightarrow ('a \Rightarrow 'b \Rightarrow bool) \Rightarrow 'b\ mem \Rightarrow bool \\
&v \triangleright_R s \\
&= (\forall m.\ cmem\ m \longrightarrow \\
&\qquad (\textbf{let}\ (v', m') = run_state\ s\ m\ \textbf{in}\ R\ v\ v' \wedge cmem\ m'))
\end{aligned}
$$

However, we still do not have a means of expressing the second part of our sentence. To this end, we use the *function relator* $(\Rightarrow\!\!\!\!\!\Rightarrow)$:

$$
\begin{aligned}
&(\Rightarrow\!\!\!\!\!\Rightarrow) :: \\
&\quad ('a \Rightarrow 'c \Rightarrow bool) \\
&\quad \Rightarrow ('b \Rightarrow 'd \Rightarrow bool) \Rightarrow ('a \Rightarrow 'b) \Rightarrow ('c \Rightarrow 'd) \Rightarrow bool \\
&R \Rightarrow\!\!\!\!\!\Rightarrow S = (\lambda f\ g.\ \forall x\ y.\ R\ x\ y \longrightarrow S\ (f\ x)\ (g\ y))
\end{aligned}
$$

Spelled out, we have $(R \Rightarrow\!\!\!\!\!\Rightarrow S)\ f\ g$ if for any values x and y that are related by R, the values $f\ x$ and $g\ y$ are related by S.

We can finally state a proof rule for application:

$$x \triangleright x_m \wedge f \triangleright_{(=)} \Rightarrow\!\!\!\!\!\Rightarrow_{(\triangleright)} f_m \longrightarrow f\ x \triangleright f_m\ .\ x_m \tag{18.4}$$

In our concrete example, we apply it once to the goal

$$fib\ (n + 1) + fib\ n \triangleright \langle\!\langle \lambda a.\ \langle\!\langle \lambda b.\ \langle\!\langle a + b \rangle\!\rangle \rangle\!\rangle\ .\ (fib_4\ (n + 1))\ .\ (fib_4\ n)$$

solve the first premise with the induction hypotheses, and arrive at

$$(+)\ (fib\ (n + 1)) \triangleright_{(=)} \Rightarrow\!\!\!\!\!\Rightarrow_{(\triangleright)} \langle\!\langle \lambda a.\ \langle\!\langle \lambda b.\ \langle\!\langle a + b \rangle\!\rangle \rangle\!\rangle\ .\ (fib_4\ (n + 1))$$

Our current rule for application (18.4) does not match this goal. Thus we need to generalize it. In addition, we need a new rule for *return*, and a rule for $(\Rightarrow\!\!\!\!\!\Rightarrow)$. To summarize, we need the following set of theorems about our consistency relation, applying them wherever they match syntactically to finish the proof of (18.2):

$$R\ x\ y \longrightarrow x \triangleright_R \langle\!\langle y \rangle\!\rangle$$

$$x \triangleright_R x_m \wedge f \triangleright_R \Rightarrow\!\!\!\!\!\Rightarrow_S f_m \longrightarrow f\ x \triangleright_S f_m\ .\ x_m$$

$$(\forall x\ y.\ R\ x\ y \longrightarrow S\ (f\ x)\ (g\ y)) \longrightarrow (R \Rightarrow\!\!\!\!\!\Rightarrow S)\ f\ g$$

The theorem we aimed for initially

$$fst\ (run_state\ (fib_4\ n)\ empty) = fib\ n \tag{18.1}$$

is now a trivial corollary of $fib\ n \triangleright fib_4\ n$. Note that by reading the equation from right to left, we have an easy way to make the memoization transparent to an end-user of *fib*.

18.3 Details of Memoization *

In this section, we will look at some further details of the memoization process and sketch how it can be applied beyond *fib*. First note that our approach of memoization hinges on two rather independent components: We transform the original program to use the state monad, to thread (an *a priori* arbitrary) state through the program. Only at the call sites of recursion, we then introduce the memoization functionality by issuing lookups and updates to the memory (as implemented by $memo_1$). We will name this first process *monadification*. For the second component, many different memory implementations can be used, as long as we can define $memo_1$ and prove its characteristic theorem (18.3). For details on this, the reader is referred to the literature. Here, we want to turn our attention towards monadification.

To discuss some of the intricacies of monadification, let us first stick with *fib* for a bit longer and consider the following alternative definition (which is mathematically equivalent but not the same program):

$$fib\ n = (\textbf{if}\ n = 0\ \textbf{then}\ 0\ \textbf{else}\ 1 + sum_list\ (map\ fib\ [0..<n - 1]))$$

We have not yet seen how to handle two ingredients of this program: constructs like *if-then-else* or case-combinators; and higher-order functions such as *map*.

It is quite clear how *if-then-else* can be lifted to the state monad:

$$if_m :: bool\ mem \Rightarrow 'a\ mem \Rightarrow 'a\ mem \Rightarrow 'a\ mem$$
$$if_m\ b_m\ x_m\ y_m = b_m \gg\!= (\lambda b.\ \textbf{if}\ b\ \textbf{then}\ x_m\ \textbf{else}\ y_m)$$

By following the structure of the terms, we can also deduce a proof rule for if_m:

$$b \triangleright b_m \wedge x \triangleright_R x_m \wedge y \triangleright_R y_m \longrightarrow$$
$$(\textbf{if}\ b\ \textbf{then}\ x\ \textbf{else}\ y) \triangleright_R if_m\ b_m\ x_m\ y_m$$

* If you are just interested in the dynamic programming algorithms of the following sections, this section can safely be skipped on first reading.

However, suppose we want to apply this proof rule to our new equation for *fib*. We will certainly need the knowledge of whether $n = 0$ to make progress in the correctness proof. Thus we make our rule more precise:

$$b \triangleright b_m \wedge (b \longrightarrow x \triangleright_R x_m) \wedge (\neg\ b \longrightarrow y \triangleright_R y_m) \longrightarrow$$
$$(\textbf{if}\ b\ \textbf{then}\ x\ \textbf{else}\ y) \triangleright_R if_m\ b_m\ x_m\ y_m$$

How can we lift *map* to the state monad level? Consider its defining equations:

$$map\ f\ [] = []$$
$$map\ f\ (x\ \#\ xs) = f\ x\ \#\ map\ f\ xs$$

We can follow the pattern we used to monadify *fib* to monadify *map*:

$$map_m{}'\ f\ [] = \langle\!\langle [] \rangle\!\rangle$$
$$map_m{}'\ f\ (x\ \#\ xs) = \langle\!\langle \lambda a.\ \langle\!\langle \lambda b.\ \langle\!\langle a\ \#\ b \rangle\!\rangle \rangle\!\rangle \rangle\!\rangle\ .\ (\langle\!\langle f \rangle\!\rangle\ .\ \langle\!\langle x \rangle\!\rangle)\ .\ (map_m{}'\ f\ xs)$$

We have obtained a function $map_m{}'$ of type

$$('a \Rightarrow 'b\ mem) \Rightarrow 'a\ list \Rightarrow 'b\ list\ mem$$

This is not yet compatible with our scheme of lifting function application to $(.)$. We need a function of type

$$(('a \Rightarrow 'b\ mem) \Rightarrow ('a\ list \Rightarrow 'b\ list\ mem)\ mem)\ mem$$

because *map* has two arguments and we need one layer of the state monad for each of its arguments. Therefore we simply define

$$map_m = \langle\!\langle \lambda f.\ \langle\!\langle map_m{}'\ f \rangle\!\rangle \rangle\!\rangle$$

For inductive proofs about the new definition of *fib*, we also need the knowledge that *fib* is recursively applied only to smaller values than n when computing *fib n*. That is, we need to know which values f is applied to in *map f xs*. We can encode this knowledge in a proof rule for *map*:

$$xs = ys \wedge (\forall x.\ x \in set\ ys \longrightarrow f\ x \triangleright_R f_m\ x) \longrightarrow$$
$$map\ f\ xs \triangleright_{list_all2\ R}\ map_m\ .\ \langle\!\langle f_m \rangle\!\rangle\ .\ \langle\!\langle ys \rangle\!\rangle$$

The relator *list_all2* lifts R to a pairwise relation on lists:

$$list_all2\ R\ xs\ ys = (|xs| = |ys| \wedge (\forall i{<}|xs|.\ R\ (xs\ !\ i)\ (ys\ !\ i)))$$

To summarize, here is a fully memoized version of the alternative definition of *fib*:

```
fib_m :: nat ⇒ nat mem
fib_m = memo_1 n
 (if_m ⟪n = 0⟫ ⟪0⟫
   (⟪λa. ⟪1 + a⟫⟫ .
    (⟪λa. ⟪sum_list a⟫⟫ . (map_m . ⟪fib_m'⟫ . ⟪[0..<n − 1]⟫)))))
```

The correctness proof for fib_m is analogous to the one for fib_4, once we have proved the new rules discussed above.

At the end of this section, we note that the techniques that were sketched above also extend to case-combinators and other higher-order functions. Most of the machinery for monadification and the corresponding correctness proofs can be automated in Isabelle [79]. Finally note that none of the techniques we used so far are specific to *fib*. The only parts that have to be adopted are the definitions of $memo_1$ and *cmem*. In Isabelle, this can be done by simply instantiating a locale.

This concludes the discussion of the fundamentals of our approach towards verified dynamic programming. We now turn to the study of two typical examples of dynamic programming algorithms: the Bellman-Ford algorithm and an algorithm for computing optimal binary search trees.

18.4 The Bellman-Ford Algorithm

Calculating shortest paths in weighted graphs is a classic algorithmic task that we all encounter in everyday situations, such as planning the fastest route to drive from A to B. In this scenario we can view streets as edges in a graph and nodes as street crossings. Every edge is associated with a weight, e.g. the time to traverse a street. We are interested in the path from A to B with minimum weight, corresponding to the fastest route in the example. Note that in this example it is safe to assume that all edge weights are non-negative.

Some applications demand negative edge weights as well. Suppose, we transport ourselves a few years into the future, where we have an electric car that can recharge itself via solar cells while driving. If we aim for the most energy-efficient route from A to B, a very sunny route could then incur a negative edge weight.

The **Bellman-Ford algorithm** is a classic dynamic programming solution to the *single-destination shortest path problem* in graphs with negative edge weights. That is, we are given a directed graph with negative edge weights

and some target vertex (known as *sink*), and we want to calculate the weight of the shortest (i.e. minimum weight) paths from every vertex to the sink. Figure 18.2 shows an example of such a graph.

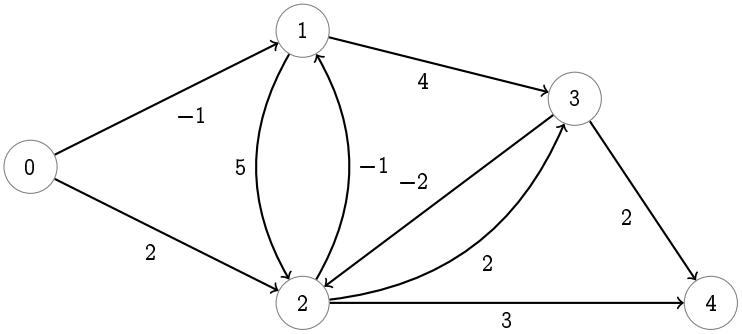


Fig. 18.2. Example of a weighted directed graph

Formally, we will take a simple view of graphs. We assume that we are given a number of nodes numbered $0,\ldots,n$, and some sink $t \in \{0..n\}$ (thus $n = t = 4$ in the example). Edge weights are given by a function $W :: int \Rightarrow int \Rightarrow int\ extended$. The type $int\ extended$ extends the natural numbers with positive and negative infinity:

datatype $'a\ extended = Fin\ 'a \mid \infty \mid -\infty$

We refrain from giving the explicit definition of addition and comparison on this domain, and rely on your intuition instead. A weight assignment $W\ i\ j = \infty$ means that there is no edge from i to j. The purpose of $-\infty$ will become clear later.

18.4.1 Deriving a Recursive Solution

The main idea of the algorithm is to consider paths in order of increasing length *in the number of edges*. In the example, we can immediately read off the weights of the shortest paths to the sink that use only one edge: only nodes 2 and 3 are directly connected to the sink, with edge weights 3 and 2, respectively; for all others the weight is infinite. How can we now calculate the minimum weight paths (to the sink) with at most two edges? For node 3, the weight of the shortest path with at most two edges is: either the weight of the path with one edge; or the weight of the edge from node 3 to node 2 plus the weight of the path with one edge from node 2 to the sink. Because $-2 +$

$3 = 1 \leq 2$, we get a new minimum weight of 1 for node 3. Following the same scheme, we can iteratively calculate the minimum path weights given in table 18.1.

i/v	0	1	2	3	4
0	∞	∞	∞	∞	0
1	∞	∞	3	2	0
2	5	6	3	1	0
3	5	5	3	1	0
4	4	5	3	1	0

Table 18.1. The minimum weights of paths from vertices $v = 0 \ldots 4$ to t that use at most $i = 0 \ldots 4$ edges.

The analysis we just ran on the example already gives us a clear intuition on all we need to deduce a dynamic program: a recursion on sub-problems, in this case to compute the weight of shortest paths with at most $i + 1$ edges from the weights of shortest paths with at most i edges. To formalize this recursion, we first define the notion of a minimum weight path from some node v to t with at most i edges, denoted as $OPT\ i\ v$:

$$
\begin{aligned}
&OPT :: nat \Rightarrow nat \Rightarrow int\ extended \\
&OPT\ i\ v \\
&= Min\ (\{weight\ (v\ \#\ xs\ @\ [t])\ |\ |xs| + 1 \leq i \wedge set\ xs \subseteq \{0..n\}\} \cup \\
&\qquad \{\textbf{if}\ t = v\ \textbf{then}\ 0\ \textbf{else}\ \infty\}) \\
\\
&weight :: (nat \Rightarrow nat \Rightarrow int\ extended) \Rightarrow nat\ list \Rightarrow int\ extended \\
&weight\ _\ [_] = 0 \\
&weight\ W\ (v\ \#\ w\ \#\ xs) = W\ v\ w + weight\ W\ (w\ \#\ xs)
\end{aligned}
$$

If $i = 0$, things are simple:

$$OPT\ 0\ v = (\textbf{if}\ t = v\ \textbf{then}\ 0\ \textbf{else}\ \infty)$$

A shortest path that constitutes $OPT\ (i + 1)\ v$ uses either at most i or exactly $i + 1$ edges. That is, $OPT\ (i + 1)\ v$ is either $OPT\ i\ v$, or the weight of the edge from v to any of its neighbours w plus $OPT\ i\ w$:

$$
\begin{aligned}
&OPT\ (i + 1)\ v \\
&= min\ (OPT\ i\ v)\ (Min\ \{W\ v\ w + OPT\ i\ w\ |\ w \leq n\})
\end{aligned}
$$

Proof. We prove this equality by proving two inequalities:

($lhs \leq rhs$) For this direction, we essentially need to show that every path on the rhs is covered by the lhs, which is trivial.

($lhs \geq rhs$) We skip the cases where $OPT\ (i + 1)\ v$ is trivially 0 or ∞ (i.e. where it is given by the singleton set in the definition of OPT). Thus consider some xs such that $OPT\ (i + 1)\ v = weight\ (v \# xs\ @\ [t])$, $|xs| \leq i$, and $set\ xs \subseteq \{0..n\}$. The cases where $|xs| < i$ or $i = 0$ are trivial. Otherwise, we have $OPT\ (i + 1)\ v = W\ v\ (hd\ xs) + weight\ (xs\ @\ [t])$ by definition of *weight*, and $OPT\ i\ (hd\ xs) \leq weight\ (xs\ @\ [t])$ by definition of OPT. Therefore, we can show:

$$OPT\ (i + 1)\ v \geq W\ v\ (hd\ xs) + OPT\ i\ (hd\ xs) \geq rhs$$

□

We can turn these equations into a recursive program:

$bf :: nat \Rightarrow nat \Rightarrow int\ extended$
$bf\ 0\ v = (\textbf{if}\ t = v\ \textbf{then}\ 0\ \textbf{else}\ \infty)$
$bf\ (i + 1)\ v$
$= min_list\ (bf\ i\ v\ \#\ map\ (\lambda w.\ W\ v\ w + bf\ i\ w)\ [0..<n + 1])$

It is obvious that we can prove correctness of bf by induction:

$$bf\ i\ v = OPT\ i\ v$$

18.4.2 Negative Cycles

Have we solved the initial problem now? The answer is "not quite" because we have ignored one additional complication. Consider our example table 18.1 again. The table stops at path length five because no shorter paths with more edges exist. For this example, five corresponds to the number of nodes, which bounds the length of the longest *simple* path. However, is it the case that we will never find shorter non-simple paths in other graphs? The answer is "no". If a graph contains a *negative reaching* cycle, i.e. a cycle with a negative sum of edge weights from which the sink is reachable, then we can use it arbitrarily often too find shorter and shorter paths.

Luckily, we can use the Bellman-Ford algorithm to detect this situation by examining the relationship of $OPT\ n$ and $OPT\ (n + 1)$. The following proposition summarizes the key insight:

> The graph contains a negative reaching cycle if and only if there exists a $v \leq n$ such that $OPT\ (n + 1)\ v < OPT\ n\ v$

Proof. If there is no negative reaching cycle, then all shortest paths are either simple or contain superfluous cycles of weight 0. Thus, we have $OPT\ (n + 1)\ v = OPT\ n\ v$ for all $v \leq n$.

Otherwise, there is a negative reaching cycle $ys = a \# xs \mathbin{@} [a]$ with $weight\ ys < 0$. Working towards a contradiction, assume that $OPT\ n\ v \leq OPT\ (n + 1)\ v$ for all $v \leq n$. Using the recursion we proved above, this implies $OPT\ n\ v \leq W\ v\ u + OPT\ n\ u$ for all u, $v \leq n$. By applying this inequality to the nodes in $a \# xs$, we can prove the inequality

$$\begin{aligned}&sum_list\ (map\ (OPT\ n)\ ys)\\ &\leq sum_list\ (map\ (OPT\ n)\ ys) + weight\ ys\end{aligned}$$

This implies $0 \leq weight\ ys$, which yields the contradiction. □

This means we can use *bf* to detect the existence of negative reaching cycles by computing one more round, i.e. $bf\ (n + 1)\ v$ for all v. If nothing changes in this step, we know that there are no negative reaching cycles and that $bf\ n$ correctly represents the shortest path weights. Otherwise, there has to be a negative reaching cycle.

Finally, we can use memoization to obtain an efficient implementation that solves the single-destination shortest path problem. Applying our memoization technique from above, we first obtain a memoizing version bf_m of *bf*. We then define the following program:

```
bellman_ford ::
  ((nat × nat, int extended) mapping, int extended list option) state
bellman_ford
= iter_bf (n, n) ≫=
    (λ_. map_m' (bf_m n) [0..<n + 1] ≫=
          (λxs. map_m' (bf_m (n + 1)) [0..<n + 1] ≫=
                (λys. ⟪if xs = ys then Some xs else None⟫)))
```

Here, $iter_bf\ (n, n)$ just computes the values from $bf_m\ 0\ 0$ to $bf_m\ n\ n$ in a row-by-row manner. Using the reasoning principles that were described above (for *fib*), we can then prove that *bellman_ford* indeed solves its intended task correctly (*shortest v* is the length of the shortest path from v to t):

$$\begin{aligned}&(\forall i{\leq}n.\ \forall j{\leq}n.\ -\infty < W\ i\ j) \longrightarrow\\ &fst\ (run_state\ bellman_ford\ empty)\\ &= (\textbf{if}\ contains_negative_reaching_cycle\ \textbf{then}\ None\\ &\quad \textbf{else}\ Some\ (map\ shortest\ [0..<n + 1]))\end{aligned}$$

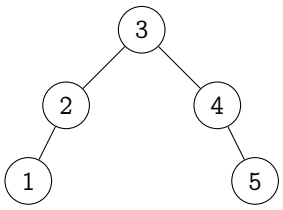
Here, *shortest* is defined analogously to OPT but for paths of unbounded length.

18.5 Optimal Binary Search Trees

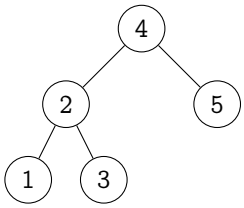
In this book, we have studied various tree data structures that guarantee logarithmic running time bounds for operations such as lookups and updates into the tree. These bounds were usually worst-case and did not take into account any information about the actual series of queries that are to be issued to the data structure. In this section, instead, we want to focus on *binary search trees* that minimize the amount of work that needs to be done when the distribution of keys in a *sequence of lookup operations* is known in advance.

More formally, we want to study the following problem. We are given a list $[i..j]$ of integers ranging from i to j and a function $p :: int \Rightarrow nat$ that maps each key in the range to a frequency with which this key is searched for. Our goal is to find a binary search tree that minimizes the expected number of comparisons when presented with a sequence of lookup operations for keys in the range $[i..j]$ that adhere to the distribution given by p.

As an example, consider the range $[1..5]$ with probabilities $[10, 30, 15, 25, 20]$. This tree



incurs an expected value of 2.15 comparison operations. However, the minimal expected value is 2 and is achieved by this tree:



Our task is equivalent to minimizing the *weighted path length* (or *cost*) as we did for Huffman encodings (Chapter 13). Recall that the weighted path length is the sum of the frequencies of every node in the tree multiplied by its depth in the tree. It fulfills the following (recursive) equations:

$$cost\ \langle\rangle = 0$$
$$cost\ \langle l, k, r\rangle$$

$$= (\sum_{k \in set_tree\ l}\ p\ k)\ +\ cost\ l\ +\ p\ k\ +\ cost\ r\ +\ (\sum_{k \in set_tree\ r}\ p\ k)$$

The difference of our task compared to finding an optimal Huffman encoding is the constraint that the resulting tree needs to be *sorted*, making it hard to deploy a similar greedy solution. Instead, we want to come up with a dynamic programming solution and thus need to find a way to subdivide the problem.

18.5.1 Deriving a Recursive Solution

The key insight into the problem is that subtrees of optimal binary search trees are also optimal. The left and right subtrees of the root must be optimal, since if we could improve either one, we would also get a better tree for the complete range of keys. This motivates the following definition:

$$\begin{aligned}
&wpl\ W\ i\ j\ \langle\rangle = 0\\
&wpl\ W\ i\ j\ \langle l,\ k,\ r\rangle\\
&= wpl\ W\ i\ (k\ -\ 1)\ l\ +\ wpl\ W\ (k\ +\ 1)\ j\ r\ +\ W\ i\ j\\
&W\ i\ j = (\sum_{k\ =\ i}^{j}\ p\ k)
\end{aligned}$$

It is easy to see that $wpl\ W\ i\ j$ is just a reformulation of $cost\ t$:

$$inorder\ t = [i..j] \longrightarrow wpl\ W\ i\ j\ t = cost\ t$$

We can actually forget about the original frequencies p and just optimize $wpl\ W\ i\ j$ for some fixed weight function $W :: int \Rightarrow int \Rightarrow nat$.

The binary search tree t that contains the keys $[i..j]$ and minimizes $wpl\ W\ i\ j\ t$ has some root k with $[i..j] = [i..k-1]\ @\ k\ \#\ [j+1..k]$. Its left and right subtrees need to be minimal again, i.e. minimize $wpl\ W\ i\ (k-1)$ and $wpl\ W\ (k+1)\ j$. This yields the following recursive functions for computing the minimal weighted path length (min_wpl) and a corresponding binary search tree (opt_bst):

```
min_wpl :: int ⇒ int ⇒ nat
min_wpl i j
= (if j < i then 0
   else min_list
        (map (λk. min_wpl i (k − 1) + min_wpl (k + 1) j +
                  W i j)
```

```
            [i..j]))

opt_bst :: int ⇒ int ⇒ int tree
opt_bst i j
= (if j < i then ⟨⟩
    else argmin (wpl W i j)
           (map (λk. ⟨opt_bst i (k − 1), k, opt_bst (k + 1) j⟩) [i..j]))
```

Here *argmin f xs* returns the rightmost $x \in set\ xs$ such that $f\ x$ is minimal among xs (i.e. $f\ x \leq f\ y$ for all $y \in set\ xs$).

To prove that *min_wpl* and *opt_bst* are correct, we want to show two properties: *min_wpl i j* should be a lower bound of *wpl W i j t* for any search tree t for $[i..j]$, and *min_wpl i j* should correspond to the weight of an actual search tree, namely *opt_bst i j*. Formally, we prove the following properties:

$$inorder\ t = [i..j] \longrightarrow min_wpl\ i\ j \leq wpl\ W\ i\ j\ t$$

$$inorder\ (opt_bst\ i\ j) = [i..j]$$

$$wpl\ W\ i\ j\ (opt_bst\ i\ j) = min_wpl\ i\ j$$

The first property is easily proved by computation induction on *wpl W*, and the second and third property are equally easily proved using computation induction on *min_wpl*.

If W is constructed from p as above, we can derive the following correctness theorems referring to the original problem:

$$inorder\ t = [i..j] \longrightarrow min_wpl\ W\ i\ j \leq cost\ t$$

$$cost\ (opt_bst\ W\ i\ j) = min_wpl\ W\ i\ j$$

18.5.2 Memoization

We can apply the memoization techniques that were discussed above to efficiently compute *min_wpl* and *opt_bst*. The only remaining caveat is that W also needs to be computed efficiently from the distribution p. If we just use the defining equality $W\ i\ j = (\sum_{k=i}^{j} p\ k)$, the computation of W is unnecessarily costly. Another way is to memoize W itself, using the following recursion:

$$W\ p\ i\ j = (\textbf{if}\ i \leq j\ \textbf{then}\ W\ p\ i\ (j-1) + p\ j\ \textbf{else}\ 0)$$

This yields a memoizing version W_m' and a theorem that connects it to W:

$$W\ p\ i\ j \triangleright W_m'\ p\ i\ j$$

We can now iterate $W_m'\ p\ i\ n$ for $i = 0 \ldots n$ to pre-compute all relevant values of $W\ p\ i\ j$:

$$W_c\ p\ n = snd\ (run_state\ (map_m'\ (\lambda i.\ W_m'\ p\ i\ n)\ [0..n])\ empty)$$

Using the correctness theorem for map_m' from above, it can easily be shown that this yields a consistent memory:

$$cmem\ (W_c\ p\ n)$$

We can show the following equation for computing W

$$W\ p\ i\ j = (\textbf{case}\ (W_c\ p\ n)\ (i, j)\ \textbf{of}\ None \Rightarrow W\ p\ i\ j \mid Some\ x \Rightarrow x)$$

Note that the *None* branch will only be triggered when indices outside of $0 \ldots n$ are accessed. Finally, we can use W_c to pass the pre-computed values of W to *opt_bst*:

```
opt_bst' :: (int ⇒ nat) ⇒ int ⇒ int ⇒ int tree
opt_bst' p i j ≡
let M = W_c p j;
    W = λi j. case M (i, j) of None ⇒ W p i j | Some x ⇒ x
in opt_bst W i j
```

18.5.3 Optimizing the Recursion

While we have applied some trickery to obtain an efficient implementation of the simple dynamic programming algorithm expressed by *opt_bst*, we still have not arrived at the solution that is currently known to be most efficient. The most efficient known algorithm to compute optimal binary search trees due to Knuth [41] is a slight variation of *opt_bst* and relies on the following observation.

Let $R\ i\ j$ denote the maximal root of any optimal binary search for $[i..j]$:

$$\begin{aligned} &R\ i\ j \\ &= argmin\ (\lambda k.\ w\ i\ j + min_wpl\ i\ (k - 1) + min_wpl\ (k + 1)\ j)\ [i..j] \end{aligned}$$

It can be shown that $R\ i\ j$ is bounded by $R\ i\ (j - 1)$ and $R\ (i + 1)\ j$:

$$i < j \longrightarrow R\ i\ (j - 1) \leq R\ i\ j \wedge R\ i\ j \leq R\ (i + 1)\ j$$

The proof of this fact is rather involved and the details can be found in the references provided at the end of this section.

With this knowledge, we can make the following optimization to *opt_bst*:

```
opt_bst₂ :: int ⇒ int ⇒ int tree

opt_bst₂ i j
= (if j < i then ⟨⟩
   else if i = j then ⟨⟨⟩, i, ⟨⟩⟩
        else let left = root (opt_bst₂ i (j − 1));
                 right = root (opt_bst₂ (i + 1) j)
             in argmin (wpl i j)
                (map (λk. ⟨opt_bst₂ i (k − 1), k,
                           opt_bst₂ (k + 1) j⟩)
                  [left..right]))
```

You may wonder whether this change really incurs an asymptotic runtime improvement. Indeed, it can be shown that it improves the algorithm's runtime by a factor of $O(n)$. For a fixed search tree size $d = i - j$, the total number of recursive computations is given by the following telescoping series:

$$
\begin{aligned}
&d \leq n \longrightarrow \\
&(\textstyle\sum j = d..n.\ \textbf{let}\ i = j - d\ \textbf{in}\ R\ (i + 1)\ j - R\ i\ (j - 1) + 1) \\
&= R\ (n - d + 1)\ n - R\ 0\ (d - 1) + n - d + 1
\end{aligned}
$$

This quantity is bounded by $2 \cdot n$, which implies that the overall number of recursive calls is bounded by $O(n^2)$.

Bibliographic Remarks

The original $O(n^2)$ algorithm for Binary Search Trees is due to Knuth [41]. Yao later explained this optimization more elegantly in his framework of "quadrilateral inequalities" [80]. Nipkow and Somogyi follow Yao's approach in their Isabelle formalization [59], on which the last subsection of this chapter is based. The other parts of this chapter are based on a paper by Wimmer *et al.* [79] and its accompanying Isabelle formalization [78]. The formalization also contains further examples of dynamic programming algorithms, including solutions for the Knapsack and the minimum edit distance problems, and the CYK algorithm.

19

Amortized Analysis

Consider a k-bit binary counter and a sequence of increment (by one) operations on it where each one starts from the least significant bit and keeps flipping the 1s until a 0 is encountered (and flipped). Thus the worst-case running time of an increment is $O(k)$ and a sequence of n increments takes time $O(nk)$. However, this analysis is very coarse: in a sequence of increments there are many much faster ones (for half of them the least significant bit is 0!). It turns out that a sequence of n increments takes time $O(n)$. Thus the average running time of each increment is $O(1)$. Amortized analysis is the analysis of the running time of a sequence of operations on some data structure by upper-bounding the average running time of each operation.

As the example of the binary counter shows, the amortized running time for a single call of an operation can be much better than the worst-case time. Thus amortized analysis is unsuitable in a real-time context where worst-case bounds on every call of an operation are required.

Below we assume that there is a single data structure (in the example: the binary counter) that is an argument to each of the operations and may also be returned as a result.

Amortized analysis is valid only if the data structure is used in a single-threaded manner. The binary counter shows why: start from 0, increment the counter until all bits are 1, then increment that counter value again and again, without destroying it. Each of those increments takes time $O(k)$ and you can do that as often as you like, thus subverting the analysis. In an imperative language you can easily avoid this "abuse" by making the data structure stateful: every operation modifies the state of the data structure. This shows that amortized analysis has an imperative flavour. In a purely functional language, monads can be used to guarantee single-threadedness.

19.1 The Potential Method

The potential method is a particular technique for amortized analysis. The key idea is to define a potential function Φ from the data structure to non-negative numbers. The potential of the data structure is like a savings account that cheap calls can pay into (by increasing the potential) to compensate for later expensive calls (which decrease the potential). In a nutshell: the less "balanced" a data structure is, the higher its potential should be because it will be needed to pay for the impending restructuring.

The **amortized running time** is defined as the actual running time plus the difference in potential, i.e. the potential after the call minus the potential before it. If the potential increases, the amortized running time is higher than the actual running time and we pay the difference into our savings account. If the potential decreases, the amortized running time is lower than the actual running time and we take something out of our savings account to pay for the difference.

More formally, we are given some data structure with operations f, g, etc on it with corresponding time functions T_f, T_g etc. We are also given a potential function Φ. Then the amortized running time function A_f for f is defined as follows:

$$A_f\ s = T_f\ s + \Phi\ (f\ s) - \Phi\ s \tag{19.1}$$

where s is the data structure under consideration; f may also have additional parameters. Given a sequence of data structure states $s_0, \ldots, s_n$ where $s_{i+1} = f_i\ s_i$, it is not hard to see that

$$\sum_{i=0}^{n-1} A_{f_i}\ s_i = \sum_{i=0}^{n-1} T_{f_i}\ s_i + \Phi\ s_n - \Phi\ s_0$$

If we assume (for simplicity) that $\Phi\ s_0 = 0$, then it follows immediately that the amortized running time of the whole sequence is an upper bound of the actual running time (because Φ is non-negative). This observation becomes useful if we can bound $A_f\ s$ by some closed term $u_f\ s$. Typical examples for $u_f\ s$ are constants, logarithms etc. Then we can conclude that f has constant, logarithmic etc amortized complexity. Thus the only proof obligation is

$$A_f\ s \leq u_f\ s$$

possibly under the additional assumption *invar* s if the data structure comes with an invariant *invar*.

In the sequel we assume that s_0 is some fixed value, typically "empty", and that its potential is 0.

How do we analyze operations that combine two data structures, e.g. the union of two sets? Their amortized complexity can be defined in analogy to (19.1):

$$A_f\ s_1\ s_2 = T_f\ s_1\ s_2 + \Phi\ (f\ s_1\ s_2) - (\Phi\ s_1 + \Phi\ s_2)$$

So far we implicitly assumed that all operations return the data structure as a result, otherwise $\Phi\ (f\ s)$ does not make sense. How should we analyze so-called **observer functions** that do not modify the data structure but return a value of some other type? Amortized analysis does not make sense here because the same observer can be applied multiple times to the same data structure value without modifying it. Classical worst-case complexity is needed, unless the observer does modify the data structure as a side effect or by returning a new value. Then one can perform an amortized analysis that ignores the returned observer value (but not the time it takes to compute it).

19.2 Examples

19.2.1 Binary Counter

The binary counter is represented by a list of booleans where the head of the list is the least significant bit. The increment operation and its running time are easily defined:

```
incr :: bool list ⇒ bool list
incr [] = [True]
incr (False # bs) = True # bs
incr (True # bs) = False # incr bs

T_incr :: bool list ⇒ real
T_incr [] = 1
T_incr (False # _) = 1
T_incr (True # bs) = T_incr bs + 1
```

The potential of a counter is the number of *True*'s because they increase T_{incr}:

```
Φ :: bool list ⇒ real
Φ bs = |filter (λx. x) bs|
```

Clearly the potential is never negative.

The amortized complexity of *incr* is 2:

$$T_{incr}\ bs + \Phi\ (incr\ bs) - \Phi\ bs = 2$$

This can be proved automatically by induction on *bs*.

Bibliographic Remarks

Amortized analysis is due to Tarjan [72]. Introductions to it can be found in most algorithm textbooks. This chapter is based on earlier work by Nipkow [53, 55] which also formalizes the meta-theory of amortized analysis.

20

Queues

20.1 Queue Specification

A **queue** can be viewed as a glorified list with function *enq* for adding an element to the end of the list and function *first* for accessing and *deq* for removing the first element. This is the full ADT:

ADT $Queue =$
interface $empty :: 'q$
$enq :: 'a \Rightarrow 'q \Rightarrow 'q$
$deq :: 'q \Rightarrow 'q$
$first :: 'q \Rightarrow 'a$
$is_empty :: 'q \Rightarrow bool$
abstraction $list :: 'q \Rightarrow 'a\ list$
invariant $invar :: 'q \Rightarrow bool$
specification $list\ empty = []$
$invar\ q \longrightarrow list\ (enq\ x\ q) = list\ q\ @\ [x]$
$invar\ q \longrightarrow list\ (deq\ q) = tl\ (list\ q)$
$invar\ q \wedge list\ q \neq [] \longrightarrow first\ q = hd\ (list\ q)$
$invar\ q \longrightarrow is_empty\ q = (list\ q = [])$
$invar\ empty$
$invar\ q \longrightarrow invar\ (enq\ x\ q)$
$invar\ q \longrightarrow invar\ (deq\ q)$

A trivial implementation is as a list, but then *enq* is linear in the length of the queue. To improve this we consider two more sophisticated implementations. First, a simple implementation where every operation has amortized constant complexity. Second, a tricky "real time" implementation where every operation has worst-case constant complexity.

20.2 Queues as Pairs of Lists

The queue is implemented as a pair of lists (fs, rs), the front and rear lists. Function *enq* adds elements to the head of the rear rs and *deq* removes elements from the head of the front fs. When fs becomes empty, it is replaced by $rev\ rs$ (and rs is emptied) — the reversal ensures that now the oldest element is at the head. Hence rs is really the reversal of the rear of the queue but we just call it the rear. The abstraction function is obvious:

```
list :: 'a list × 'a list ⇒ 'a list
list (fs, rs) = fs @ rev rs
```

Clearly *enq* and *deq* are constant-time until the front becomes empty. Then we need to reverse the rear which takes linear time (if it is implemented by *itrev*, see Section 1.5.1). But we can pay for this linear cost up front by paying a constant amount for each call of *enq*. Thus we arrive at amortized constant time. See below for the formal treatment.

The implementation is shown in Figure 20.1. Of course $empty = ([], [])$. Function *norm* performs the reversal of the rear once the front becomes

```
norm :: 'a list × 'a list ⇒ 'a list × 'a list
norm (fs, rs) = (if fs = [] then (itrev rs [], []) else (fs, rs))

enq :: 'a ⇒ 'a list × 'a list ⇒ 'a list × 'a list
enq a (fs, rs) = norm (fs, a # rs)

deq :: 'a list × 'a list ⇒ 'a list × 'a list
deq (fs, rs) = (if fs = [] then (fs, rs) else norm (tl fs, rs))

first :: 'a list × 'a list ⇒ 'a
first (a # _, _) = a

is_empty :: 'a list × 'a list ⇒ bool
is_empty (fs, _) = (fs = [])
```

Fig. 20.1. Queue as a pair of lists

empty. Why does not only *deq* but also *enq* call *norm*? Because otherwise $enq\ x_n\ (...(enq\ x_1\ empty)...)$ would result in $([], [x_n, ..., x_1])$ and *first* would

become an expensive operation because it would requires the reversal of the rear. Thus we need to avoid queues $([], rs)$ where $rs \neq []$. Thus *norm* guarantees the following invariant:

```
invar :: 'a list × 'a list ⇒ bool
invar (fs, rs) = (fs = [] ⟶ rs = [])
```

Functional correctness, i.e. proofs of the properties in the ADT *Queue*, are straightforward. Let us now turn to the amortized running time analysis. The time functions are shown in Appendix B.6.

For the amortized analysis we define the potential function

```
Φ :: 'a list × 'a list ⇒ nat
Φ (_, rs) = |rs|
```

because $|rs|$ is the amount we have accumulated by charging 1 for each *enq*. This is enough to pay for the eventual reversal. Now it is easy to prove that both *enq* and *deq* have amortized constant running time:

$$T_{enq}\ a\ (fs, rs) + \Phi\ (enq\ a\ (fs, rs)) - \Phi\ (fs, rs) \leq 4$$
$$T_{deq}\ (fs, rs) + \Phi\ (deq\ (fs, rs)) - \Phi\ (fs, rs) \leq 3$$

The two observer functions *first* and *is_empty* have constant running time.

20.3 A Real Time Implementation

This sections presents the **Hood-Melville queue**, a tricky implementation that improves upon the representation in the previous Section by preemptively performing reversals over a number of operations before they are required.

20.3.1 Stepped Reversal

Breaking down a reversal operation into multiple steps can be done using the following function:

```
rev_step :: 'a list × 'a list ⇒ 'a list × 'a list
rev_step (x # xs, ys) = (xs, x # ys)
rev_step ([], ys) = ([], ys)
```

where $x \# xs$ is the list being reversed, and $x \# ys$ is the partial reversal result. Thus, to reverse a list of size 3 one should call *rev_step* 3 times:

rev_step ([1, 2, 3], []) = ([2, 3], [1])
rev_step (*rev_step* ([1, 2, 3], [])) = ([3], [2, 1])
rev_step (*rev_step* (*rev_step* ([1, 2, 3], []))) = ([], [3, 2, 1])

Note that each call to *rev_step* takes constant time since its definition is non-recursive.

Using the notation f^n for the n-fold composition of function f we can state a simple inductive lemma:

Lemma 20.1. $rev_step^{|xs|}\ (xs,\ ys) = ([],\ rev\ xs\ @\ ys)$

As a special case this implies $rev_step^{|xs|}\ (xs,\ []) = ([],\ rev\ xs)$.

20.3.2 A Real Time Intuition

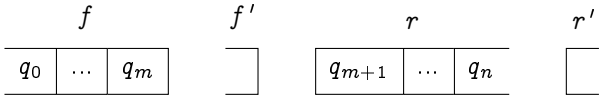
Hood-Melville queues are similar to those presented in Section 20.2 in that they use a pair of lists $(f,\ r)$ (front and rear — for succinctness we drop the s's now) to achieve constant running time *deq* and *enq*. However, they avoid a costly reversal operation once f becomes empty by preemptively computing a new front $fr = f\ @\ rev\ r$ one step at a time using *rev_step* as enqueueing and dequeueing operations occur. The process that generates fr consists of three phases:

1. Reverse r to form r', which is the tail end of fr
2. Reverse f to form f'
3. Reverse f' onto r' to form fr

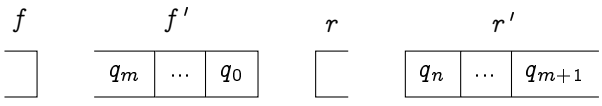
All three phases can be described in terms of *rev_step* as follows:

1. $r' = snd\ (rev_step^{|r|}\ (r,\ []))$
2. $f' = snd\ (rev_step^{|f|}\ (f,\ []))$
3. $fr = snd\ (rev_step^{|f'|}\ (f',\ r'))$

Phases (1) and (2) are independent and can be performed at the same time, hence, when starting from this configuration



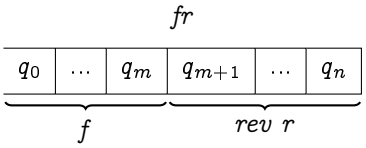
after $max\ |f|\ |r|$ steps of reversal the state would be the following:



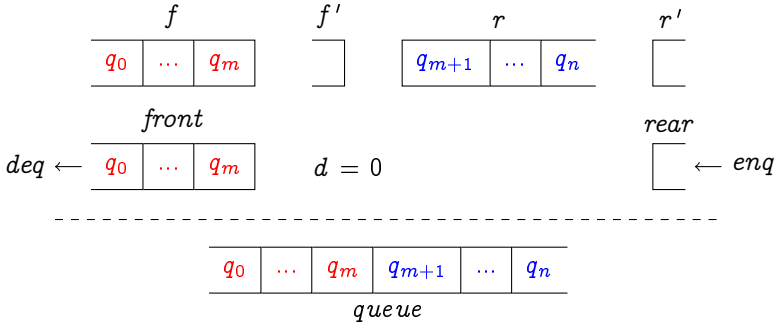
Phase (3) reverses f' onto r' fr to obtain the same result as a call to *list*:

$$\begin{aligned}
fr &= snd\ (rev_step^{|f'|}\ (f', r')) && \text{by definition of } fr \\
&= rev\ f' \mathbin{@} r' && \text{using Lemma 20.1} \\
&= rev\ f' \mathbin{@} snd\ (rev_step^{|r|}\ (r, [])) && \text{by definition of } r' \\
&= rev\ f' \mathbin{@} rev\ r && \text{using Lemma 20.1} \\
&= rev\ (snd\ (rev_step^{|f|}\ (f, []))) \mathbin{@} rev\ r' && \text{by definition of } f' \\
&= rev\ (rev\ f) \mathbin{@} rev\ r && \text{using Lemma 20.1} \\
&= f \mathbin{@} rev\ r && \text{by } rev \text{ involution}
\end{aligned}$$

The resulting front list fr contains all elements previously in f and r:



A Hood-Melville queue spreads all reversal steps across queue altering operations requiring careful bookkeeping. To achieve this gradual reversal, additional lists *front* and *rear* are used for enqueuing and dequeuing, while internal operations rely only on f, f', r, and r'. At the start of the reversal process *rear* is copied into r and emptied; similarly, *front* is copied into f, but its contents are kept as they might need to be dequeued. Moreover, to avoid using elements from f or f' that may have been removed from *front*, a counter d records the number of dequeuing operations that have occurred since the reversal process started; this way, only $|f'| - d$ elements are appended into r to form fr. Once the reversal finishes fr become the new *front* and the internal lists are cleared. When the queue is not being reversed all operations are performed in a manner similar to previous implementations. The configuration of a queue at the beginning of the reversal process is as follows:



20.3.3 The Reversal Strategy

A crucial detail of this implementation is determining at which point the reversal process should occur. The strategy is to start once $|rear|$ becomes larger than $|front|$, and ensure that all reversal steps are done before *front* runs out of elements or *rear* becomes larger than the new front (fr).

With this strategy, once $|rear| = n+1$ and $|front| = n$, the reversal processes starts. The first two phases take $n + 1$ steps ($max\ |front|\ |rear|$) to generate f' and r', and the third phase produces fr in n steps. A complete reversal takes $2 \cdot n + 1$ steps. Because the queue can only perform n *deq* operations before *front* is exhausted, $2 \cdot n + 1$ steps must be performed in at most n operations. This can be achieved by performing the first two steps in the operation that causes *rear* to become larger than *front* and two more steps in each subsequent operation. Therefore, $2 \cdot (n + 1)$ steps can occur before *front* is emptied, allowing the reversal process to finish in time.

Finally, since at most n *enq* or *deq* operations can occur during reversal, the largest possible *rear* has length n (only *enq* ops), while the smallest possible fr has length $n + 1$ (only *deq* ops). Thus, after the reversing process has finished the new front (fr) is always larger than *rear*.

20.3.4 Implementation

Queues are implemented using the following record type:

```
record 'a queue = lenf   :: nat
                  front  :: 'a list
                  status :: 'a status
                  rear   :: 'a list
                  lenr   :: nat
```

In a nutshell, a record is a product type with named fields and "built-in" construction, selection, and update operations. Values of *'a queue* are constructed using *make* :: *nat* ⇒ *'a list* ⇒ *'a status* ⇒ *'a list* ⇒ *nat* ⇒ *'a queue* were each argument corresponds to one of the fields of the record in canonical order. Additionally, given a queue q we can obtain the value in field *front* with *front q*, and update its content using $q(\!|front := []|\!)$. Multiple updates can be composed as $q(\!|front := [],\ rear := []|\!)$.

All values in the queue along with its internal state are stored in the various fields of *'a queue*. Fields *front* and *rear* contain the lists over which all queue operations are performed. The length of *front* and *rear* is recorded in *lenf* and *lenr* (respectively) to avoid calling *length* whose complexity is

not constant. Finally, *status* tracks the current reversal phase of the queue in a *'a status* value.

```
datatype 'a status =
  Idle |
  Rev nat ('a list) ('a list) ('a list) |
  App nat ('a list) ('a list) |
  Done
```

Each value of *'a status* represents either a phase of reversal or the queue's normal operation. Constructor *Idle* signals that no reversal is being performed. *Rev ok f f' r r'* corresponds to phases (1) and (2) where the lists *f*, *f'*, *r*, and *r'* are used for the reversal steps of the front and the rear. The *App ok f' r'* case corresponds to phase (3) where both lists are appended to form the new front (*fr*). In both *App* and *Rev*, the first argument *ok* (of type *nat*) keeps track of the number of elements in *f'* that have not been removed from the queue, effectively $ok = |f'| - d$, where *d* is the number of *deq* operations that have occurred so far. Lastly, *Done fr* marks the end of the reversal process and contains only the new front list *fr*.

In the implementation, all of the steps of reversal operations in the queue are performed by functions *exec* and *invalidate*; they ensure at each step that the front list being computed is kept consistent w.r.t. the contents and operations in the queue.

Function *exec* :: *'a status* ⇒ *'a status* performs the incremental reversal of the front list by altering the queue's *status* one step at a time in accordance with the reversal phases. Following the strategy described in Section 20.3.3, all queue operations call *exec* twice to be able to finish the reversal in time. On *Idle* queues *exec* has no effect. The implementation of *exec* is an extension of *rev_step* with specific considerations for each *status* value and is defined as follows:

```
exec :: 'a status ⇒ 'a status
exec (Rev ok (x # f) f' (y # r) r')
= Rev (ok + 1) f (x # f') r (y # r')
exec (Rev ok [] f' [y] r') = App ok f' (y # r')
exec (App 0 f' r') = Done r'
exec (App ok (x # f') r') = App (ok - 1) f' (x # r')
exec s = s
```

If the *status* is *Rev ok f f' r r'*, then *exec* performs two (or one if $f = []$) simultaneous reversal steps from *f* and *r* into *f'* and *r'*; moreover *ok* is

incremented if a new element has been added to f'. Once f is exhausted and r is a singleton list, the remaining element is moved into r' and the *status* is updated to the next phase of reversal. In the $App\ ok\ f'\ r'$ phase, *exec* moves elements from f' to r' until $ok = 0$, at which point r' becomes the new front by transitioning into $Done\ r'$. In all other cases *exec* behaves as the identity function. As is apparent from its implementation, a number of assumptions are required for *exec* to function properly and eventually produce *Done*. These assumption are discussed in Section 20.3.5.

If an element is removed from the queue during the reversal process, it also needs to be removed from the new front list (fr) being computed. Function *invalidate* is used to achieve this:

```
invalidate :: 'a status ⇒ 'a status
invalidate (Rev ok f f' r r') = Rev (ok − 1) f f' r r'
invalidate (App 0 f' (_ # r')) = Done r'
invalidate (App ok f' r') = App (ok − 1) f' r'
invalidate s = s
```

By decreasing the value of *ok*, the number of elements from f' that are moved into r' in phase (3) is reduced, since *exec* might produce *Done* early, once $ok = 0$, ignoring the remaining elements of f'. Furthermore, since f' is a reversal of the front list, elements left behind in its tail correspond directly to those being removed from the queue.

The rest of the implementation is shown below. Auxiliary function *exec2*, as its name suggests, applies *exec* twice and updates the queue accordingly if *Done* is returned.

```
exec2 :: 'a queue ⇒ 'a queue
exec2 q = (case exec (exec q) of
             Done fr ⇒ q(|status = Idle, front = fr|) |
             newstatus ⇒ q(|status = newstatus|))

check :: 'a queue ⇒ 'a queue
check q
= (if lenr q ≤ lenf q then exec2 q
   else let newstate = Rev 0 (front q) [] (rear q) []
        in exec2
           (q(|lenf := lenf q + lenr q, status := newstate, rear := [],
               lenr := 0|)))
```

```
empty :: 'a queue
empty = make 0 [] Idle [] 0

first :: 'a queue ⇒ 'a
first q = hd (front q)

enq :: 'a ⇒ 'a queue ⇒ 'a queue
enq x q = check (q(|rear := x # rear q, lenr := lenr q + 1|))

deq :: 'a queue ⇒ 'a queue
deq q
= check
    (q(|lenf := lenf q − 1, front := tl (front q),
        status := invalidate (status q)|))
```

The two main queue operations, *enq* and *deq*, alter *front* and *rear* as expected, with additional updates to *lenf* and *lenr* to keep track of the their length. To perform all "internal" operations, both functions call *check*. Additionally, *deq* uses *invalidate* to mark elements as removed.

Function *check* calls *exec2* if *lenr* is not larger than *lenf*. Otherwise a reversal process is initiated: *rear* is emptied and *lenr* is set to 0; *lenf* is increased to the size of the whole queue since, conceptually, all element are now in the soon-to-be-computed front; the status *newstate* is initialized as described at the beginning of Section 20.3.2.

The time complexity of this implementation is clearly constant, since there are no recursive functions.

20.3.5 Functional Correctness

To show this implementation is an instance of the ADT *Queue*, we need a number of invariants to ensure the consistency of *'a queue* values are preserved by all operations.

Initially, as hinted by the definition of *exec*, values of type *'a status* should have specific properties to guarantee a *Done* result after a (finite) number of calls to *exec*. The predicate *inv_st* defines these properties as follows:

```
inv_st :: 'a status ⇒ bool
inv_st (Rev ok f f' r r') = (|f| + 1 = |r| ∧ |f'| = |r'| ∧ ok ≤ |f'|)
```

```
inv_st (App ok f' r') = (ok ≤ |f'| ∧ |f'| < |r'|)
inv_st Idle = True
inv_st (Done _) = True
```

First, *inv_st* ensures for pattern, *Rev ok f f' r r'* that arguments f and r follow the reversal strategy, and counter *ok* is only ever increased as elements are added to f'. Similarly, for *App ok f' r'*, it must follow that r' remains larger than f', and $|f'|$ provides an upper bound for *ok*. All other patterns trivially fulfill the invariant.

The consistency of *'a queue* values is then defined as an extension of *inv_st* and considers all the other fields in the queue. The predicate *invar* defines the required properties:

```
invar q
= (lenf q = |front_list q| ∧ lenr q = |rear_list q| ∧
   lenr q ≤ lenf q ∧
   (case status q of
    Rev ok f f' _ _ ⇒
      2 · lenr q ≤ |f'| ∧ ok ≠ 0 ∧
      2 · |f| + ok + 2 ≤ 2 · |front q|
    | App ok _ r ⇒ 2 · lenr q ≤ |r| ∧ ok + 1 ≤ 2 · |front q|
    | _ ⇒ True) ∧
   (∃ rest. front_list q = front q @ rest) ∧
   (∄ fr. status q = Done fr) ∧
   inv_st (status q))
```

Predicate *lenr q* = |*rear_list q*| ensures *lenr* is equal to the length of the queue's rear, where function *rear_list q*, defined as (*rev* ∘ *rear*) *q*, produces the rear list in canonical order. Likewise, *lenf q* = |*front_list q*| matches *lenf* to the queue's front. However, function *front_list* warrants special attention as it must compute the list representing the front of the queue even during a reversal. *front_list* is defined as follows:

```
front_list :: 'a queue ⇒ 'a list
front_list q
= (case status q of Idle ⇒ front q
   | Rev ok f f' r r' ⇒ rev (take ok f') @ f @ rev r @ r'
   | App ok f' x ⇒ rev (take ok f') @ x | Done f ⇒ f)
```

For case $App\ ok\ f'\ r'$, the front list corresponds to the final result of the stepped reversal (20.1) but only elements in f' that are still in the queue, denoted by $take\ ok\ f'$, are considered. Analogously for $Rev\ ok\ f\ f'\ r\ r'$, both stepped reversal results are appended and only relevant elements in f' are used, however, rear lists r and r' are reversed again to achieve canonical order.

Continuing with *invar*, inequality $lenr\ q \leq lenf\ q$ is the main invariant in our reversal strategy, and by the previous two equalities must holds even as internal operations occur. Furthermore, predicate $\exists\ rest.\ front_list\ q = front\ q\ @\ rest$ ensures $front\ q$ is contained within $front_list\ q$, thus preventing any mismatch between the internal state and the queue's front. Given that *exec*2 is the only function that manipulates a queue's *status*, it holds that $\nexists\ fr.\ status\ q = Done\ fr$ since any internal *Done* result is replaced by *Idle*.

The case distinction on $status\ q$ places size bounds on internal lists *front* and *rear* ensuring the front does not run out of elements and the rear never grows beyond $lenr\ q \leq lenf\ q$. In order to clarify some of formulations used in this part of invar, consider the following correspondences, which hold once the reversal process starts:

- $lenr\ q$ corresponds to the number of *enq* operations performed so far, and $2 \cdot lenr\ q$ denotes the *exec* applications in those operations.
- $|front\ q|$ corresponds to the number of *deq* operations that can be performed before $front\ q$ is exhausted. Therefore, $2 \cdot |front\ q|$ is the minimum number of *exec* applications the queue will be able to do at any given point.
- On $Rev\ ok\ f\ f'\ r\ r'$ status, $|f'|$ corresponds to the number of applications of *exec* performed so far and the internal length of front being constructed. Expression $|r|$ is the analogous for a $App\ ok\ f\ r$.
- From a well formed $App\ ok\ f\ r$ it takes $ok + 1$ applications of *exec* to reach *Done*. Since, the base case of *App* is obtained after ok applications, and the transition into *Done* takes an extra step.
- From a well formed $Rev\ ok\ f\ f'\ r\ r'$ it takes $2 \cdot |f'| + ok + 2$ applications of *exec* to reach *Done*. Since, the base case of *Rev* is obtained after $|f'|$ applications (incrementing ok by the same amount), the transitioning into *App* takes one step, and $ok + |f'|$ extra steps are need to reach *Done* from *App*.

In the $Rev\ ok\ f\ f'\ r\ r'$ case, $2 \cdot lenr\ q \leq |f'|$ ensures f' grows larger with every *enq* operation and the internal list is at least twice the length of the queue's rear. Additionally, the value of ok cannot be 0 as this either marks the beginning of a reversal which calls *exec*2 immediately, or signals that elements in $front\ q$ have run out. Finally, to guarantee the reversal process can finish before the $front\ q$ is exhausted the number of *exec* applications

before reaching *Done* must be less than the minimum number of applications possible, denoted by $2 \cdot |f| + ok + 2 \leq 2 \cdot |front\ q|$.

Case *App ok f r* has similar invariants, with equation $2 \cdot lenr\ q \leq |r|$ bounding the growth of r as it was previously done with f'. Moreover, $ok + 1 \leq 2 \cdot |front\ q|$ ensures *fron q* is not exhausted before the reversal is completed.

Using *invar* we can prove the properties of *enq* and *deq* required by the *Queue* ADT using this abstraction following abstraction function:

$$list :: 'a\ queue \Rightarrow 'a\ list$$
$$list\ q = front_list\ q\ @\ rear_list\ q$$

Proofs follow mostly by case distinction on the *status* field. Invariants *invar* and *inv_st* are used to filter invalid *'a queue* values, and to simplify terms based on their relationship with other terms (e.g. $f = [] \wedge |f| + 1 = |r| \longrightarrow (\exists x.\ r = [x])$). All other properties in *Queue* are straightforward.

Bibliographic Remarks

The representation of queues as pairs of lists is due to Burton [11]. The Hood-Melville queues are due to Hood and Melville [30]. The implementation is based on the presentation by Okasaki [62].

21

Splay Trees

Splay trees are fascinating self-organizing search trees. Self-organizing means that the tree structure is modified upon access (including *isin* queries) to improve the performance of subsequent operations. Concretely, every splay tree operation moves the element concerned to the root. Thus splay trees excel in applications where a small fraction of the entries are the targets of most of the operations. In general, splay trees perform as well as any static binary search tree.

Splay trees have two drawbacks. First, their performance guarantees (logarithmic running time of each operation) are only amortized. Self-organizing does not mean self-balancing: splay trees can become unbalanced, in contrast to, for example, red-black trees. Second, because *isin* modifies the tree, splay trees are less convenient to use in a purely functional language.

21.1 Implementation

The central operation on splay trees is the *splay* function shown in Figure 21.1. It rotates the given element x to the root of the tree if x is already in the tree. Otherwise the last element found before the search for x hits a leaf is rotated to the root.

Function *isin* has a trivial implementation in terms of *splay*:

```
isin :: 'a tree ⇒ 'a ⇒ bool
isin t x = (case splay x t of ⟨⟩ ⇒ False | ⟨_, a, _⟩ ⇒ x = a)
```

Except that *splay* creates a new tree that needs to be returned from a proper *isin* as well to achieve the amortized logarithmic complexity (see the discussion of observer functions at the end of Section 19.1). This is why splay trees

```
splay x ⟨AB, b, CD⟩
= (case cmp x b of
    LT ⇒ case AB of
           ⟨⟩ ⇒ ⟨AB, b, CD⟩ |
           ⟨A, a, B⟩ ⇒
             case cmp x a of
             LT ⇒ if A = ⟨⟩ then ⟨A, a, ⟨B, b, CD⟩⟩
                  else case splay x A of
                         ⟨A₁, a′, A₂⟩ ⇒ ⟨A₁, a′, ⟨A₂, a, ⟨B, b, CD⟩⟩⟩ |
             EQ ⇒ ⟨A, a, ⟨B, b, CD⟩⟩ |
             GT ⇒ if B = ⟨⟩ then ⟨A, a, ⟨B, b, CD⟩⟩
                  else case splay x B of
                         ⟨B₁, b′, B₂⟩ ⇒ ⟨⟨A, a, B₁⟩, b′, ⟨B₂, b, CD⟩⟩ |
    EQ ⇒ ⟨AB, b, CD⟩ |
    GT ⇒ case CD of
           ⟨⟩ ⇒ ⟨AB, b, CD⟩ |
           ⟨C, c, D⟩ ⇒
             case cmp x c of
             LT ⇒ if C = ⟨⟩ then ⟨⟨AB, b, C⟩, c, D⟩
                  else case splay x C of
                         ⟨C₁, c′, C₂⟩ ⇒ ⟨⟨AB, b, C₁⟩, c′, ⟨C₂, c, D⟩⟩ |
             EQ ⇒ ⟨⟨AB, b, C⟩, c, D⟩ |
             GT ⇒ if D = ⟨⟩ then ⟨⟨AB, b, C⟩, c, D⟩
                  else case splay x D of
                         ⟨D₁, d, D₂⟩ ⇒ ⟨⟨⟨AB, b, C⟩, c, D₁⟩, d, D₂⟩)
```

Fig. 21.1. Function *splay*

are inconvenient in functional languages. For the moment we ignore this aspect and stick with the above *isin* because it has the type required by the *Set* ADT.

The implementation of *insert* x t in Figure 21.2 is straightforward: let $\langle l, a, r\rangle$ = *splay* x t; if $a = x$, return $\langle l, a, r\rangle$; otherwise make x the root of a suitable recombination of l, a and r.

The implementation of *delete* x t in Figure 21.3 starts similarly: let $\langle l, a, r\rangle$ = *splay* x t; if $a \neq x$, return $\langle l, a, r\rangle$. Otherwise follow the deletion-by-replacing paradigm (Section 5.2.1): if $l \neq \langle\rangle$, splay the maximal element m in l to the root and replace x with it. Note that *splay_max* returns a tree that is just a glorified pair: if $t \neq \langle\rangle$ then *splay_max* t is of the form $\langle t', m, \langle\rangle\rangle$. The definition *splay_max* $\langle\rangle = \langle\rangle$ is not really needed (*splay_max* is always called with non-$\langle\rangle$ argument) but some lemmas can be stated more slickly with this definition.

```
insert :: 'a ⇒ 'a tree ⇒ 'a tree
insert x t
= (if t = ⟨⟩ then ⟨⟨⟩, x, ⟨⟩⟩
    else case splay x t of
         ⟨l, a, r⟩ ⇒ case cmp x a of
                       LT ⇒ ⟨l, x, ⟨⟨⟩, a, r⟩⟩ |
                       EQ ⇒ ⟨l, a, r⟩ |
                       GT ⇒ ⟨⟨l, a, ⟨⟩⟩, x, r⟩)
```

Fig. 21.2. Function *insert*

```
delete :: 'a ⇒ 'a tree ⇒ 'a tree
delete x t
= (if t = ⟨⟩ then ⟨⟩
    else case splay x t of
         ⟨l, a, r⟩ ⇒
           if x ≠ a then ⟨l, a, r⟩
           else if l = ⟨⟩ then r
                else case splay_max l of ⟨l', m, _⟩ ⇒ ⟨l', m, r⟩)

splay_max :: 'a tree ⇒ 'a tree
splay_max ⟨⟩ = ⟨⟩
splay_max ⟨A, a, ⟨⟩⟩ = ⟨A, a, ⟨⟩⟩
splay_max ⟨A, a, ⟨B, b, CD⟩⟩
= (if CD = ⟨⟩ then ⟨⟨A, a, B⟩, b, ⟨⟩⟩
    else case splay_max CD of ⟨C, c, D⟩ ⇒ ⟨⟨⟨A, a, B⟩, b, C⟩, c, D⟩)
```

Fig. 21.3. Functions *delete* and *splay_max*

21.2 Correctness

The *inorder* approach of Section 5.4 applies. Because the details are a bit different (everything is reduced to *splay*) we present the top-level structure.

The following easy inductive properties are used implicitly in a number of subsequent proofs:

$$splay\ a\ t = \langle\rangle \longleftrightarrow t = \langle\rangle$$

$$splay_max\ t = \langle\rangle \longleftrightarrow t = \langle\rangle$$

Correctness of *isin*

$$sorted\ (inorder\ t) \longrightarrow isin\ t\ x = (x \in set\ (inorder\ t))$$

follows directly from this easy inductive property of *splay*:

$$splay\ x\ t = \langle l, a, r\rangle \wedge sorted\ (inorder\ t) \longrightarrow (x \in set\ (inorder\ t)) = (x = a)$$

Correctness of *insert* and *delete*

$$sorted\ (inorder\ t) \longrightarrow inorder\ (insert\ x\ t) = ins_list\ x\ (inorder\ t)$$
$$sorted\ (inorder\ t) \longrightarrow inorder\ (delete\ x\ t) = del_list\ x\ (inorder\ t)$$

relies on the following characteristic inductive properties of *splay*:

$$inorder\ (splay\ x\ t) = inorder\ t \tag{21.1}$$
$$sorted\ (inorder\ t) \wedge splay\ x\ t = \langle l, a, r\rangle \longrightarrow sorted\ (inorder\ l\ @\ x\ \#\ inorder\ r)$$

Correctness of *delete* also needs the inductive property

$$splay_max\ t = \langle l, a, r\rangle \wedge sorted\ (inorder\ t) \longrightarrow inorder\ l\ @\ [a] = inorder\ t \wedge r = \langle\rangle$$

Note that $inorder\ (splay\ x\ t) = inorder\ t$ is also necessary to justify the proper *isin* that returns the newly created tree as well.

Automation of the above proofs requires the lemmas in Figure 5.2 together with a few additional lemmas about *sorted*, *ins_list* and *del_list* that can be found in the Isabelle proofs.

Recall from Section 5.4 that correctness of *insert* and *delete* implies that they preserve $bst = sorted \circ inorder$. Similarly, (21.1) implies that *splay* preserves *bst*. Thus we may assume the invariant *bst* in the amortized analysis.

These two easy size lemmas are used implicitly below:

$$|splay\ a\ t| = |t| \qquad |splay_max\ t| = |t|$$

21.3 Amortized Analysis

This section shows that *splay*, insertion and deletion all have amortized logarithmic complexity.

We define the potential Φ of a tree as the sum of the potentials φ of all nodes:

$$\Phi :: 'a\ tree \Rightarrow real$$
$$\Phi\ \langle\rangle = 0$$
$$\Phi\ \langle l, a, r\rangle = \varphi\ \langle l, a, r\rangle + \Phi\ l + \Phi\ r$$

$$\varphi\ t \equiv \log_2\ |t|_1$$

The central result is the amortized complexity of *splay*. Function T_{splay} is shown in Appendix B.7. We follow (19.1) and define

$$A_{splay}\ a\ t = T_{splay}\ a\ t + \Phi\ (splay\ a\ t) - \Phi\ t$$

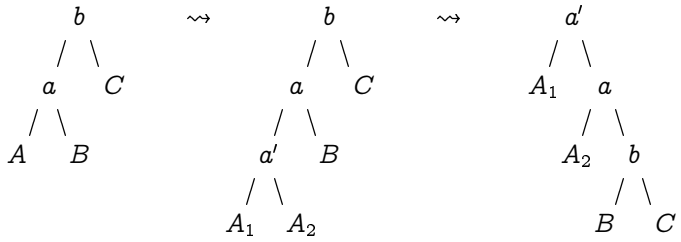
First we consider the case where the element is in the tree:

Theorem 21.1. $bst\ t \wedge \langle l, x, r\rangle \in subtrees\ t \longrightarrow$
$A_{splay}\ x\ t \le 3 \cdot (\varphi\ t - \varphi\ \langle l, x, r\rangle) + 1$

Proof by induction on the computation of *splay*. The base cases involving $\langle\rangle$ are impossible. For example, consider the call *splay* $x\ t$ where $t = \langle\langle\rangle, b, C\rangle$ and $x < b$: from $\langle l, x, r\rangle \in subtrees\ t$ it follows that $x \in set_tree\ t$ but because $bst\ t$ and $x < b$ this implies that $x \in set_tree\ \langle\rangle$, a contradiction. There are three feasible base cases. The case $t = \langle _, x, _\rangle$ is easy. We consider one of the two other symmetric cases. Let $t = \langle\langle A, x, B\rangle, b, C\rangle$ and $t' = splay\ x\ t = \langle A, x, \langle B, b, C\rangle\rangle$.

$$\begin{aligned}
&A_{splay}\ x\ t = \Phi\ t' - \Phi\ t + 1 && \text{by definition of } A_{splay} \text{ and } T_{splay}\\
&= \varphi\ t' + \varphi\ \langle B, b, C\rangle - \varphi\ t - \varphi\ \langle A, x, B\rangle + 1 && \text{by definition of } \Phi\\
&= \varphi\ \langle B, b, C\rangle - \varphi\ \langle A, x, B\rangle + 1 && \text{by definition of } \varphi\\
&\le \varphi\ t - \varphi\ \langle A, x, B\rangle + 1 && \text{because } \varphi\ \langle B, b, C\rangle \le \varphi\ t\\
&\le 3 \cdot (\varphi\ t - \varphi\ \langle A, x, B\rangle) + 1 && \text{because } \varphi\ \langle A, x, B\rangle \le \varphi\ t\\
&= 3 \cdot (\varphi\ t - \varphi\ \langle l, x, r\rangle) + 1 && \text{because } bst\ t \wedge \langle l, x, r\rangle \in subtrees\ t
\end{aligned}$$

There are four inductive cases. We consider two of them, the other two are symmetric variants. First the so-called zig-zig case:



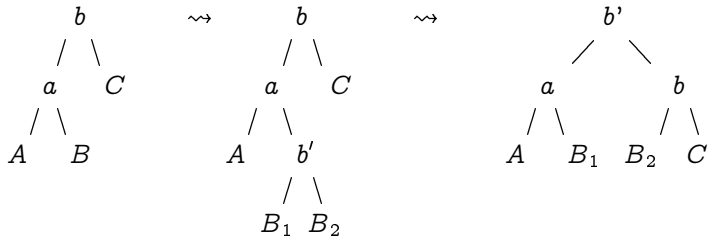
This is the case where $x < a < b$ and $A \neq \langle\rangle$. On the left we have the input and on the right the output of *splay* x. Because $A \neq \langle\rangle$, *splay* $x\ A = \langle A_1, a', A_2\rangle =: A'$ for some A_1, a' and A_2. The intermediate tree is obtained by replacing A by A'. This tree is shown for illustration purpose only; in the algorithm the right tree is constructed directly from the left one. Let $X = \langle l, x, r\rangle$. Clearly $X \in subtrees\ A$. We abbreviate compound trees like $\langle A, a, B\rangle$ by the names of their subtrees, in this case AB. First note that

$$\varphi\ A_1A_2BC = \varphi\ ABC \qquad (*)$$

because $|A'| = |\mathit{splay}\ x\ A| = |A|$. We can now prove the claim:

$$
\begin{aligned}
&A_{\mathit{splay}}\ x\ ABC = T_{\mathit{splay}}\ x\ A + 1 + \Phi\ A_1A_2BC - \Phi\ ABC \\
&= T_{\mathit{splay}}\ x\ A + 1 + \Phi\ A_1 + \Phi\ A_2 + \varphi\ A_2BC + \varphi\ BC - \Phi\ A - \varphi\ AB \\
&\qquad\qquad \text{by } (*) \text{ and definition of } \Phi \\
&= T_{\mathit{splay}}\ x\ A + \Phi\ A' - \varphi\ A' - \Phi\ A + \varphi\ A_2BC + \varphi\ BC - \varphi\ AB + 1 \\
&= A_{\mathit{splay}}\ x\ A + \varphi\ A_2BC + \varphi\ BC - \varphi\ AB - \varphi\ A' + 1 \\
&\le 3 \cdot \varphi\ A + \varphi\ A_2BC + \varphi\ BC - \varphi\ AB - \varphi\ A' - 3 \cdot \varphi\ X + 2 \\
&\qquad\qquad \text{by IH and } X \in \mathit{subtrees}\ A \\
&= 2 \cdot \varphi\ A + \varphi\ A_2BC + \varphi\ BC - \varphi\ AB - 3 \cdot \varphi\ X + 2 \\
&\qquad\qquad \text{because } \varphi\ A = \varphi\ A' \\
&< \varphi\ A + \varphi\ A_2BC + \varphi\ BC - 3 \cdot \varphi\ X + 2 \qquad \text{because } \varphi\ A < \varphi\ AB \\
&< \varphi\ A_2BC + 2 \cdot \varphi\ ABC - 3 \cdot \varphi\ X + 1 \\
&\qquad\qquad \text{because } 1 + \lg x + \lg y < 2 \cdot \lg (x + y) \text{ if } x, y > 0 \\
&< 3 \cdot (\varphi\ ABC - \varphi\ X) + 1 \qquad\qquad \text{because } \varphi\ A_2BC < \varphi\ ABC
\end{aligned}
$$

Now we consider the so-called zig-zag case:



This is the case where $a < x < b$ and $B \neq \langle\rangle$. On the left we have the input and on the right the output of *splay* x. Because $B \neq \langle\rangle$, $\mathit{splay}\ x\ B = \langle B_1,\ b',\ B_2\rangle =: B'$ for some B_1, b' and B_2. The intermediate tree is obtained by replacing B by B'. Let $X = \langle l,\ x,\ r\rangle$. Clearly $X \in \mathit{subtrees}\ B$. The proof is very similar to the zig-zig case, the same naming conventions apply and we omit some details:

$$
\begin{aligned}
&A_{\mathit{splay}}\ x\ ABC = T_{\mathit{splay}}\ x\ A + 1 + \Phi\ AB_1B_2C - \Phi\ ABC \\
&= A_{\mathit{splay}}\ x\ B + \varphi\ AB_1 + \varphi\ B_2C - \varphi\ AB - \varphi\ B' + 1 \\
&\qquad\qquad \text{using } \varphi\ AB_1B_2C = \varphi\ ABC \\
&\le 3 \cdot \varphi\ B + \varphi\ AB_1 + \varphi\ B_2C - \varphi\ AB - \varphi\ B' - 3 \cdot \varphi\ X + 2 \\
&\qquad\qquad \text{by IH and } X \in \mathit{subtrees}\ B \\
&= 2 \cdot \varphi\ B + \varphi\ AB_1 + \varphi\ B_2C - \varphi\ AB - 3 \cdot \varphi\ X + 2 \\
&\qquad\qquad \text{because } \varphi\ B = \varphi\ B' \\
&< \varphi\ B + \varphi\ AB_1 + \varphi\ B_2C - 3 \cdot \varphi\ X + 2 \qquad \text{because } \varphi\ B < \varphi\ AB \\
&< \varphi\ B + 2 \cdot \varphi\ ABC - 3 \cdot \varphi\ X + 1 \\
&\qquad\qquad \text{because } 1 + \lg x + \lg y < 2 \cdot \lg (x + y) \text{ if } x, y > 0 \\
&< 3 \cdot (\varphi\ ABC - \varphi\ X) + 1 \qquad\qquad \text{because } \varphi\ B < \varphi\ ABC \qquad \square
\end{aligned}
$$

Because $\varphi\ \langle l,\ x,\ r\rangle \geq 1$, the above theorem implies

Corollary 21.2. $bst\ t \wedge x \in set_tree\ t \longrightarrow A_{splay}\ x\ t \leq 3 \cdot (\varphi\ t - 1) + 1$

If x is not in the tree we show that there is a y in the tree such that splaying with y would produce the same tree in the same time:

Lemma 21.3. $t \neq \langle\rangle \wedge bst\ t \longrightarrow$
$(\exists\, y{\in}set_tree\ t.\ splay\ y\ t = splay\ x\ t \wedge T_{splay}\ y\ t = T_{splay}\ x\ t)$

Element y is the last element in the tree that the search for x encounters before it hits a leaf. Naturally, the proof is by induction on the computation of *splay*.

Combining this lemma with Corollary 21.2 yields the final unconditional amortized complexity of *splay* on BSTs:

Corollary 21.4. $bst\ t \longrightarrow A_{splay}\ x\ t \leq 3 \cdot \varphi\ t + 1$

The "$-\ 1$" has disappeared to accommodate the case $t = \langle\rangle$.

The amortized analysis of insertion is straightforward now. From the amortized complexity of *splay* it follows that

Lemma 21.5. $bst\ t \longrightarrow T_{insert}\ x\ t + \Phi\ (insert\ x\ t) - \Phi\ t \leq 4 \cdot \varphi\ t + 3$

We omit the proof which is largely an exercise in simple algebraic manipulations.

The amortized analysis of deletion is similar but a bit more complicated because of the additional function *splay_max* whose amortized running time is defined as usual:

$$A_{splay_max}\ t = T_{splay_max}\ t + \Phi\ (splay_max\ t) - \Phi\ t$$

An inductive proof and then a simple case analysis yield

$t \neq \langle\rangle \longrightarrow A_{splay_max}\ t \leq 3 \cdot (\varphi\ t - 1) + 1$

$A_{splay_max}\ t \leq 3 \cdot \varphi\ t + 1$

Based on the canonical definitions of T_{delete} and T_{splay_max} the amortized logarithmic complexity of *delete* follows:

$bst\ t \longrightarrow T_{delete}\ a\ t + \Phi\ (delete\ a\ t) - \Phi\ t \leq 6 \cdot \varphi\ t + 3$

A running time analysis of *isin* is trivial because *isin* is just *splay* followed by a constant-time test.

21.4 Exercises

Exercise 21.1. Find a sequence of numbers $n_1,\ n_2,\ \ldots n_k$ such that the insertion of theses numbers one by one creates a tree of height k.

Bibliographic Remarks

Splay trees were invented and analyzed by Sleator and Tarjan [70] for which they received the 1999 ACM Paris Kanellakis Theory and Practice Award [38]. In addition to the amortized complexity as shown above they proved that splay trees perform as well as static BSTs (the Static Optimality Theorem) and conjectured that, roughly speaking, they even perform as well as any other BST-based algorithm. This Dynamic Optimality Conjecture is still open.

This chapter is based on earlier publications [69, 53, 54, 55].

22

Skew Heaps

Skew heaps are heaps in the sense of Section 14.1 and implement mergeable priority queues. Skew heaps can be viewed as a self-adjusting form of leftist heaps that attempts to maintain balance by unconditionally swapping all nodes in the merge path when merging two heaps.

22.1 Implementation of ADT *Priority_Queue_Merge*

The central operation is *merge*:

$$
\begin{aligned}
&merge :: 'a\ tree \Rightarrow 'a\ tree \Rightarrow 'a\ tree\\
&merge\ \langle\rangle\ t = t\\
&merge\ t\ \langle\rangle = t\\
&merge\ (\langle l_1,\ a_1,\ r_1\rangle =: t_1)\ (\langle l_2,\ a_2,\ r_2\rangle =: t_2)\\
&= (\textbf{if}\ a_1 \le a_2\ \textbf{then}\ \langle merge\ t_2\ r_1,\ a_1,\ l_1\rangle\ \textbf{else}\ \langle merge\ t_1\ r_2,\ a_2,\ l_2\rangle)
\end{aligned}
$$

The remaining operations ($\{\}$, *insert*, *get_min* and *del_min*) are defined as in Section 14.1.

The following properties of *merge* have easy inductive proofs:

$$|merge\ t_1\ t_2| = |t_1| + |t_2|$$

$$mset_tree\ (merge\ t_1\ t_2) = mset_tree\ t_1 + mset_tree\ t_2$$

$$heap\ t_1 \wedge heap\ t_2 \longrightarrow heap\ (merge\ t_1\ t_2)$$

Now it is straightforward to prove the correctness of the implementation w.r.t. the ADT *Priority_Queue_Merge*.

Skew heaps attempt to maintain balance, but this does not always work:

Exercise 22.1. Find a sequence of numbers $n_1, n_2, \ldots n_k$ such that the insertion of theses numbers one by one creates a tree of height k. Prove that this sequence will produce a tree of height k.

Nevertheless, insertion and deletion have amortized logarithmic complexity.

22.2 Amortized Analysis

The key is the definition of the potential. It counts the number of **right-heavy** (rh) nodes:

```
Φ :: 'a tree ⇒ int
Φ ⟨⟩ = 0
Φ ⟨l, _, r⟩ = Φ l + Φ r + rh l r

rh :: 'a tree ⇒ 'a tree ⇒ nat
rh l r = (if |l| < |r| then 1 else 0)
```

The rough intuition: because *merge* descends along the right spine, the more right-heavy nodes a tree contains, the longer *merge* takes.

Two auxiliary functions count the number of right-heavy nodes on the left spine (lrh) and left-heavy (= not right-heavy) nodes on the right spine (rlh):

```
lrh :: 'a tree ⇒ nat
lrh ⟨⟩ = 0
lrh ⟨l, _, r⟩ = rh l r + lrh l

rlh :: 'a tree ⇒ nat
rlh ⟨⟩ = 0
rlh ⟨l, _, r⟩ = 1 − rh l r + rlh r
```

The following properties have automatic inductive proofs:

$$2^{lrh\ t} \le |t| + 1 \qquad 2^{rlh\ t} \le |t| + 1$$

They imply

$$lrh\ t \le \lg |t|_1 \qquad rlh\ t \le \lg |t|_1 \tag{22.1}$$

Now we are ready for the amortized analysis. All canonical time functions can be found in Appendix B.8. The key lemma is an upper bound of the amortized complexity of *merge* in terms of lrh and rlh:

Lemma 22.1. $T_{merge}\ t_1\ t_2 + \Phi\ (merge\ t_1\ t_2) - \Phi\ t_1 - \Phi\ t_2$
$\leq lrh\ (merge\ t_1\ t_2) + rlh\ t_1 + rlh\ t_2 + 1$

Proof by induction on the computation of *merge*. We consider only the node-node case: let $t_1 = \langle l_1, a_1, r_1\rangle$ and $t_2 = \langle l_2, a_2, r_2\rangle$. W.l.o.g. assume $a_1 \leq a_2$. Let $m = merge\ t_2\ r_1$.

$$
\begin{aligned}
&T_{merge}\ t_1\ t_2 + \Phi\ (merge\ t_1\ t_2) - \Phi\ t_1 - \Phi\ t_2 \\
&= T_{merge}\ t_2\ r_1 + 1 + \Phi\ m + \Phi\ l_1 + rh\ m\ l_1 - \Phi\ t_1 - \Phi\ t_2 \\
&= T_{merge}\ t_2\ r_1 + 1 + \Phi\ m + rh\ m\ l_1 - \Phi\ r_1 - rh\ l_1\ r_1 - \Phi\ t_2 \\
&\leq lrh\ m + rlh\ t_2 + rlh\ r_1 + rh\ m\ l_1 + 2 - rh\ l_1\ r_1 && \text{by IH} \\
&= lrh\ m + rlh\ t_2 + rlh\ t_1 + rh\ m\ l_1 + 1 \\
&= lrh\ (merge\ t_1\ t_2) + rlh\ t_1 + rlh\ t_2 + 1 && \square
\end{aligned}
$$

As a consequence we can prove the following logarithmic upper bound on the amortized complexity of *merge*:

$$
\begin{aligned}
&T_{merge}\ t_1\ t_2 + \Phi\ (merge\ t_1\ t_2) - \Phi\ t_1 - \Phi\ t_2 \\
&\leq lrh\ (merge\ t_1\ t_2) + rlh\ t_1 + rlh\ t_2 + 1 && \text{by Lemma 22.1} \\
&\leq \lg |merge\ t_1\ t_2|_1 + \lg |t_1|_1 + \lg |t_2|_1 + 1 && \text{by (22.1)} \\
&\leq \lg (|t_1|_1 + |t_2|_1 - 1) + \lg |t_1|_1 + \lg |t_2|_1 + 1 \\
& && \text{because } |merge\ t_1\ t_2| = |t_1| + |t_2| \\
&\leq \lg (|t_1|_1 + |t_2|_1) + 2 \cdot \lg (|t_1|_1 + |t_2|_1) + 1 \\
& && \text{because } \lg x + \lg y \leq 2 \cdot \lg (x + y) \text{ if } x,y > 0 \\
&= 3 \cdot \lg (|t_1|_1 + |t_2|_1) + 1
\end{aligned}
$$

The amortized complexity of insertion and deletion follows easily from the complexity of *merge*:

$$T_{insert}\ a\ t + \Phi\ (insert\ a\ t) - \Phi\ t \leq 3 \cdot \lg (|t|_1 + 2) + 2$$
$$T_{del_min}\ t + \Phi\ (del_min\ t) - \Phi\ t \leq 3 \cdot \lg (|t|_1 + 2) + 2$$

Bibliographic Remarks

Skew heaps were invented by Sleator and Tarjan [71] as one of the first self-organizing data structures. Their presentation was imperative. Our presentation follows earlier work by Nipkow [53, 55] based on the functional account by Kaldewaij and Schoenmakers [37].

23

Pairing Heaps

The pairing heap is another form of a self-adjusting priority queue. Section 23.1 presents an intuitive version of pairing heaps based on lists. In the rest of the chapter we change to a slightly different presentation that leads to a more succinct amortized analysis.

23.1 Implementation via Lists

A **pairing heap** is a heap in the sense that it is a tree with the minimal element at the root — except that it is not a binary tree but a tree where each node has a list of children:

```
datatype 'a heap = Empty | Hp 'a ('a heap list)
```

The abstraction function to multisets and the invariant follow the heap paradigm:

```
mset_heap :: 'a heap ⇒ 'a multiset
mset_heap Empty = {}
mset_heap (Hp x hs) = {x} + ∑# (mset (map mset_heap hs))

pheap :: 'a heap ⇒ bool
pheap Empty = True
pheap (Hp x hs)
= (∀ h∈set hs. (∀ y∈#mset_heap h. x ≤ y) ∧ pheap h)
```

Note that *pheap* is not the full invariant. Moreover, *Empty* does not occur inside a non-empty heap.

The implementations of *empty* and *get_min* are obvious, and *insert* follows the standard heap paradigm:

```
empty = Empty

get_min :: 'a heap ⇒ 'a
get_min (Hp x _) = x

insert :: 'a ⇒ 'a heap ⇒ 'a heap
insert x h = merge (Hp x []) h
```

Function *merge* is not recursive (as in binary heaps) but simply adds one of the two heaps to the front of the top-level heaps of the other, depending on the root value:

```
merge :: 'a heap ⇒ 'a heap ⇒ 'a heap
merge h Empty = h
merge Empty h = h
merge (Hp x hsx =: hx) (Hp y hsy =: hy)
= (if x < y then Hp x (hy # hsx) else Hp y (hx # hsy))
```

Thus *merge* and *insert* have constant running time. All the work is offloaded on *del_min* which just calls *merge_pairs*:

```
del_min :: 'a heap ⇒ 'a heap
del_min Empty = Empty
del_min (Hp _ hs) = merge_pairs hs

merge_pairs :: 'a heap list ⇒ 'a heap
merge_pairs [] = Empty
merge_pairs [h] = h
merge_pairs (h1 # h2 # hs)
= merge (merge h1 h2) (merge_pairs hs)
```

Function *merge_pairs* is a compact way of expressing a two pass algorithm: on the first pass from left to right, it merges pairs of adjacent heaps (hence "pairing heap") and on the second pass it merges the results in a cascade from right to left. By reformulating the definition in terms of these two passes, we obtain a more readable formulation with the same running time:

```
del_min :: 'a heap ⇒ 'a heap
del_min Empty = Empty
del_min (Hp _ hs) = pass₂ (pass₁ hs)

pass₁ :: 'a heap list ⇒ 'a heap list
pass₁ (h₁ # h₂ # hs) = merge h₁ h₂ # pass₁ hs
pass₁ hs = hs

pass₂ :: 'a heap list ⇒ 'a heap
pass₂ [] = Empty
pass₂ (h # hs) = merge h (pass₂ hs)
```

The proof of $pass_2\ (pass_1\ hs) = merge_pairs\ hs$ is an easy induction.

Clearly *del_min* can take linear time but it will turn out that the constant-time *insert* saves enough to guarantee amortized logarithmic complexity for both insertion and deletion.

We base the correctness proofs on the *merge_pairs* version of *del_min*. From the following lemmas (all proofs are routine inductions)

$h \neq Empty \longrightarrow get_min\ h \in_\# mset_heap\ h$

$h \neq Empty \wedge pheap\ h \wedge x \in_\# mset_heap\ h \longrightarrow get_min\ h \leq x$

$mset_heap\ (merge\ h_1\ h_2) = mset_heap\ h_1 + mset_heap\ h_2$

$mset_heap\ (merge_pairs\ hs) = \sum_\#\ (image_mset\ mset_heap\ (mset\ hs))$

$h \neq Empty \longrightarrow mset_heap\ (del_min\ h) = mset_heap\ h - \{\!\{get_min\ h\}\!\}$

$pheap\ h_1 \wedge pheap\ h_2 \longrightarrow pheap\ (merge\ h_1\ h_2)$

$(\forall\, h \in set\ hs.\ pheap\ h) \longrightarrow pheap\ (merge_pairs\ hs)$

$pheap\ h \longrightarrow pheap\ (del_min\ h)$

the properties in the specifications *Priority_Queue*(*_Merge*) follow easily.

23.2 Amortized Analysis

The amortized analysis of pairing heaps is slightly simplified if we replace the above type of heaps by trees as follows: a heap $Hp\ x\ hs$ is expressed as the tree $\langle hs,\ x,\ \langle\rangle\rangle$ and a list of heaps $[Hp\ x_1\ hs_1,\ Hp\ x_2\ hs_2,\ \ldots]$ is expressed

```
empty = ⟨⟩

get_min :: 'a tree ⇒ 'a
get_min ⟨_, x, _⟩ = x

link :: 'a tree ⇒ 'a tree
link ⟨hsx, x, ⟨hsy, y, hs⟩⟩
= (if x < y then ⟨⟨hsy, y, hsx⟩, x, hs⟩ else ⟨⟨hsx, x, hsy⟩, y, hs⟩)
link hp = hp

pass₁ :: 'a tree ⇒ 'a tree
pass₁ ⟨hsx, x, ⟨hsy, y, hs⟩⟩ = link ⟨hsx, x, ⟨hsy, y, pass₁ hs⟩⟩
pass₁ hp = hp

pass₂ :: 'a tree ⇒ 'a tree
pass₂ ⟨hsx, x, hs⟩ = link ⟨hsx, x, pass₂ hs⟩
pass₂ ⟨⟩ = ⟨⟩

get_min :: 'a tree ⇒ 'a
get_min ⟨_, x, _⟩ = x

merge :: 'a tree ⇒ 'a tree ⇒ 'a tree
merge ⟨⟩ hp = hp
merge hp ⟨⟩ = hp
merge ⟨hsx, x, ⟨⟩⟩ ⟨hsy, y, ⟨⟩⟩ = link ⟨hsx, x, ⟨hsy, y, ⟨⟩⟩⟩

insert :: 'a ⇒ 'a tree ⇒ 'a tree
insert x hp = merge ⟨⟨⟩, x, ⟨⟩⟩ hp
```

Fig. 23.1. Pairing heaps via trees

as the tree $\langle hs_1, x_1, \langle hs_2, x_2, \ldots\rangle\ldots\rangle$. This simplifies the analysis because we now have to deal only with a single type, trees.

The code for the tree representation of pairing heaps is shown in Figure 23.1. We work with the $pass_1/pass_2$ version of *del_min*. The correctness proof is very similar to what we saw in the previous section. We merely display the two invariants:

```
is_root :: 'a tree ⇒ bool
is_root hp = (case hp of ⟨⟩ ⇒ True | ⟨_, _, r⟩ ⇒ r = ⟨⟩)
```

$$pheap :: 'a\ tree \Rightarrow bool$$
$$pheap\ \langle\rangle = True$$
$$pheap\ \langle l,\ x,\ r\rangle = ((\forall\, y{\in}set_tree\ l.\ x \leq y) \wedge pheap\ l \wedge pheap\ r)$$

Now we turn to the amortized analysis. The potential of a tree is the sum of the logarithms of the sizes of the subtrees:

$$\Phi :: 'a\ tree \Rightarrow real$$
$$\Phi\ \langle\rangle = 0$$
$$\Phi\ \langle l,\ x,\ r\rangle = \lg\ |\langle l,\ x,\ r\rangle| + \Phi\ l + \Phi\ r$$

These easy inductive size properties are frequently used implicitly below:

$$|link\ hp| = |hp|$$
$$|pass_1\ hp| = |hp|$$
$$|pass_2\ hp| = |hp|$$
$$is_root\ h_1 \wedge is_root\ h_2 \longrightarrow |merge\ h_1\ h_2| = |h_1| + |h_2|$$

23.2.1 Potential Differences

We can now analyze the differences in potential caused by all the queue operations. In a separate step we will derive their amortized complexities.

For insertion, the following upper bound follows trivially from the definitions:

Lemma 23.1. $is_root\ hp \longrightarrow \Phi\ (insert\ x\ hp) - \Phi\ hp \leq \lg\ (|hp| + 1)$

For *merge* it needs a bit more work:

Lemma 23.2. $h_1 = \langle hs_1,\ x_1,\ \langle\rangle\rangle \wedge h_2 = \langle hs_2,\ x_2,\ \langle\rangle\rangle \longrightarrow$
$\Phi\ (merge\ h_1\ h_2) - \Phi\ h_1 - \Phi\ h_2 \leq \lg\ (|h_1| + |h_2|) + 1$

Proof. From

$$\begin{aligned}
&\Phi\ (merge\ h_1\ h_2)\\
&= \Phi\ (link\ \langle hs_1,\ x_1,\ h_2\rangle)\\
&= \Phi\ hs_1 + \Phi\ hs_2 + \lg\ (|hs_1| + |hs_2| + 1) + \lg\ (|hs_1| + |hs_2| + 2)\\
&= \Phi\ hs_1 + \Phi\ hs_2 + \lg\ (|hs_1| + |hs_2| + 1) + \lg\ (|h_1| + |h_2|)
\end{aligned}$$

it follows that

$$\begin{aligned}
&\Phi\ (merge\ h_1\ h_2) - \Phi\ h_1 - \Phi\ h_2\\
&= \lg\ (|hs_1| + |hs_2| + 1) + \lg\ (|h_1| + |h_2|)\\
&\quad - \lg\ (|hs_1| + 1) - \lg\ (|hs_2| + 1)
\end{aligned}$$

$\leq \lg(|h_1| + |h_2|) + 1$
because $\lg(1 + x + y) \leq 1 + \lg(1 + x) + \lg(1 + y)$ if $x,y \geq 0$
□

Now we come to the core of the proof, the analysis of *del_min*. Its running time is linear in the number of nodes reachable by descending to the right (starting from the left child of the root). We denote this metric by *len*:

```
len :: 'a tree ⇒ nat
len ⟨⟩ = 0
len ⟨_, _, r⟩ = 1 + len r
```

Therefore we have to show that the potential change compensates for this linear work. Our main goal is this:

Theorem 23.3. $\Phi(del_min\ \langle hs, x, \langle\rangle\rangle) - \Phi\ \langle hs, x, \langle\rangle\rangle$
$\leq 2 \cdot \lg(|hs| + 1) - len\ hs + 2$

It will be proved in two steps: First we show that $pass_1$ frees enough potential to compensate for the work linear in *len hs* and increases the potential only by a logarithmic term. Then we show that the increase due to $pass_2$ is also only at most logarithmic. Combining these results one easily shows that the amortized running time of *del_min* is indeed logarithmic.

First we analyze the potential difference caused by $pass_1$:

Lemma 23.4. $\Phi(pass_1\ hs) - \Phi\ hs \leq 2 \cdot \lg(|hs| + 1) - len\ hs + 2$

Proof by induction on the computation of $pass_1$. The base cases are trivial. We focus on the induction step. Let $t = \langle hs_1, x, \langle hs_2, y, hs\rangle\rangle$, $n_1 = |hs_1|$, $n_2 = |hs_2|$ and $m = |hs|$.

$\Phi(pass_1\ t) - \Phi\ t$
$= \lg(n_1 + n_2 + 1) - \lg(n_2 + m + 1) + \Phi(pass_1\ hs) - \Phi\ hs$
$\leq \lg(n_1 + n_2 + 1) - \lg(n_2 + m + 1)$
$\quad + 2 \cdot \lg(m + 1) - len\ hs + 2$ by IH
$\leq 2 \cdot \lg(n_1 + n_2 + m + 1) - \lg(n_2 + m + 1) + \lg(m + 1)$
$\quad - len\ hs$ because $\lg x + \lg y + 2 \leq 2 \cdot \lg(x + y)$ if $x,y > 0$
$\leq 2 \cdot \lg(n_1 + n_2 + m + 2) - len\ hs$
$= 2 \cdot \lg|t| - len\ t + 2$
$\leq 2 \cdot \lg(|t| + 1) - len\ t + 2$ □

Now we turn to $pass_2$:

Lemma 23.5. $hs \neq \langle\rangle \longrightarrow \Phi(pass_2\ hs) - \Phi\ hs \leq \lg|hs|$

Proof by induction on hs. The base cases are trivial. The induction step (for $\langle hs_1, x, hs\rangle$) is trivial if $hs = \langle\rangle$. Assume $hs = \langle hs_2, y, r\rangle$. Now we need one more property of $pass_2$:

$$\exists\, hs_3\ z.\ pass_2\ \langle hs_2, y, r\rangle = \langle hs_3, z, \langle\rangle\rangle$$

The proof is a straightforward induction on r. This implies $|hs_3| + 1 = |hs|$ and thus

$$\begin{aligned}&\Phi\ (link\ \langle hs_1, x, pass_2\ hs\rangle) - \Phi\ hs_1 - \Phi\ (pass_2\ hs)\\ &= \lg\ (|hs_1| + |hs| + 1) + \lg\ (|hs_1| + |hs|) - \lg\ |hs| \qquad (*)\end{aligned}$$

Thus the overall claim follows:

$$\begin{aligned}&\Phi\ (pass_2\ \langle hs_1, x, hs\rangle) - \Phi\ \langle hs_1, x, hs\rangle\\ &= \Phi\ (link\ \langle hs_1, x, pass_2\ hs\rangle) - \Phi\ hs_1 - \Phi\ hs - \lg\ (|hs_1| + |hs| + 1)\\ &= \Phi\ (pass_2\ hs) - \Phi\ hs + \lg\ (|hs_1| + |hs|) - \lg\ |hs| && \text{by } (*)\\ &\leq \lg\ (|hs_1| + |hs|) && \text{by IH}\\ &\leq \lg\ |\langle hs_1, x, hs\rangle| && \square\end{aligned}$$

Corollary 23.6. $\Phi\ (pass_2\ hs) - \Phi\ hs \leq \lg\ (|hs| + 1)$

Finally we can prove Theorem 23.3:

$$\begin{aligned}&\Phi\ (del_min\ \langle hs, x, \langle\rangle\rangle) - \Phi\ \langle hs, x, \langle\rangle\rangle\\ &= \Phi\ (pass_2\ (pass_1\ hs)) - \lg\ (|hs| + 1) - \Phi\ hs\\ &\leq \Phi\ (pass_1\ hs) - \Phi\ hs && \text{by Corollary 23.6}\\ &\leq 2 \cdot \lg\ (|hs| + 1) - len\ hs + 2 && \text{by Lemma 23.4}\end{aligned}$$

23.2.2 Amortized Running Times

The canonical running time functions are displayed in Appendix B.9. It is now straightforward to derive these amortized running times:

$$is_root\ h \longrightarrow T_{insert}\ a\ h + \Phi\ (insert\ a\ h) - \Phi\ h \leq \lg\ (|h| + 1) + 1$$

$$\begin{aligned}&is_root\ h_1 \wedge is_root\ h_2 \longrightarrow\\ &T_{merge}\ h_1\ h_2 + \Phi\ (merge\ h_1\ h_2) - \Phi\ h_1 - \Phi\ h_2\\ &\leq \lg\ (|h_1| + |h_2| + 1) + 2\end{aligned}$$

They follow from the corresponding Lemmas 23.1 and 23.2.

Combining this inductive upper bound for the running time of the two passes

$$T_{pass2}\ (pass_1\ hs_1) + T_{pass1}\ hs_1 \leq len\ hs_1 + 2$$

with Theorem 23.3 yields the third and final amortized running time:

$$\begin{aligned}&is_root\ h \longrightarrow\\ &T_{del_min}\ h + \Phi\ (del_min\ h) - \Phi\ h \leq 2 \cdot \lg\ (|h| + 1) + 5\end{aligned}$$

Thus we could prove that insertion, merging and deletion all have amortized logarithmic running times.

Bibliographic Remarks

Pairing heaps were invented by Fredman *et al.* [21] as a simpler but competitive alternative to Fibonacci heaps. The authors gave the amortized analysis presented above and conjectured that it can be improved. Later research confirmed this [33, 65, 34] but the final analysis is still open. An empirical study [47] showed that pairing heaps do indeed outperform Fibonacci heaps in practice. This chapter is based on an article by Nipkow and Brinkop [55].

Part V

Appendix

A

List Library

The following functions on lists are predefined:

$$length :: 'a\ list \Rightarrow nat$$
$$|[]| = 0$$
$$|x\ \#\ xs| = |xs| + 1$$

$$(@) :: 'a\ list \Rightarrow 'a\ list \Rightarrow 'a\ list$$
$$[]\ @\ ys = ys$$
$$(x\ \#\ xs)\ @\ ys = x\ \#\ xs\ @\ ys$$

$$set :: 'a\ list \Rightarrow 'a\ set$$
$$set\ [] = \{\}$$
$$set\ (x\ \#\ xs) = \{x\} \cup set\ xs$$

$$map :: ('a \Rightarrow 'b) \Rightarrow 'a\ list \Rightarrow 'b\ list$$
$$map\ f\ [] = []$$
$$map\ f\ (x\ \#\ xs) = f\ x\ \#\ map\ f\ xs$$

$$filter :: ('a \Rightarrow bool) \Rightarrow 'a\ list \Rightarrow 'a\ list$$
$$filter\ p\ [] = []$$
$$filter\ p\ (x\ \#\ xs) = (\textbf{if}\ p\ x\ \textbf{then}\ x\ \#\ filter\ p\ xs\ \textbf{else}\ filter\ p\ xs)$$

$$concat :: 'a\ list\ list \Rightarrow 'a\ list$$
$$concat\ [] = []$$
$$concat\ (x\ \#\ xs) = x\ @\ concat\ xs$$

```
take :: nat ⇒ 'a list ⇒ 'a list
take _ [] = []
take n (x # xs) = (case n of 0 ⇒ [] | m + 1 ⇒ x # take m xs)

drop :: nat ⇒ 'a list ⇒ 'a list
drop _ [] = []
drop n (x # xs) = (case n of 0 ⇒ x # xs | m + 1 ⇒ drop m xs)

hd :: 'a list ⇒ 'a
hd (x # xs) = x

tl :: 'a list ⇒ 'a list
tl [] = []
tl (x # xs) = xs

butlast :: 'a list ⇒ 'a list
butlast [] = []
butlast (x # xs) = (if xs = [] then [] else x # butlast xs)

rev :: 'a list ⇒ 'a list
rev [] = []
rev (x # xs) = rev xs @ [x]

(!) :: 'a list ⇒ nat ⇒ 'a
(x # xs) ! n = (case n of 0 ⇒ x | k + 1 ⇒ xs ! k)

list_update :: 'a list ⇒ nat ⇒ 'a ⇒ 'a list
[][_ := _] = []
(x # xs)[i := v] = (case i of 0 ⇒ v # xs | j + 1 ⇒ x # xs[j := v])

upt :: nat ⇒ nat ⇒ nat list
[_..<0] = []
[i..<j + 1] = (if i ≤ j then [i..<j] @ [j] else [])

replicate :: nat ⇒ 'a ⇒ 'a list
replicate 0 _ = []
replicate (n + 1) x = x # replicate n x
```

```
sum_list :: 'a list ⇒ 'a
sum_list [] = 0
sum_list (x # xs) = x + sum_list xs

min_list :: 'a list ⇒ 'a
min_list (x # xs)
= (case xs of [] ⇒ x | _ # _ ⇒ min x (min_list xs))

sorted_wrt :: ('a ⇒ 'a ⇒ bool) ⇒ 'a list ⇒ bool
sorted_wrt P [] = True
sorted_wrt P (x # ys) = ((∀ y∈set ys. P x y) ∧ sorted_wrt P ys)
```

B

Time Functions

B.1 Lists

$$
\begin{aligned}
&T_{length} :: 'a\ list \Rightarrow nat \\
&T_{length}\ [] = 1 \\
&T_{length}\ (_ \ \#\ xs) = T_{length}\ xs + 1 \\
&T_{map} :: ('a \Rightarrow nat) \Rightarrow 'a\ list \Rightarrow nat \\
&T_{map}\ _\ [] = 1 \\
&T_{map}\ T_f\ (x \ \#\ xs) = T_f\ x + T_{map}\ T_f\ xs + 1 \\
&T_{filter} :: ('a \Rightarrow nat) \Rightarrow 'a\ list \Rightarrow nat \\
&T_{filter}\ _\ [] = 1 \\
&T_{filter}\ T_p\ (x \ \#\ xs) = T_p\ x + T_{filter}\ T_p\ xs + 1 \\
&T_{take} :: nat \Rightarrow 'a\ list \Rightarrow nat \\
&T_{take}\ _\ [] = 1 \\
&T_{take}\ n\ (_ \ \#\ xs) = (\textsf{case}\ n\ \textsf{of}\ 0 \Rightarrow 1 \mid n' + 1 \Rightarrow T_{take}\ n'\ xs + 1) \\
&T_{drop} :: nat \Rightarrow 'a\ list \Rightarrow nat \\
&T_{drop}\ _\ [] = 1 \\
&T_{drop}\ n\ (_ \ \#\ xs) = (\textsf{case}\ n\ \textsf{of}\ 0 \Rightarrow 1 \mid n' + 1 \Rightarrow T_{drop}\ n'\ xs + 1)
\end{aligned}
$$

$$
\begin{aligned}
&T_{length}\ xs = |xs| + 1 \\
&T_{map}\ T_f\ xs = (\textstyle\sum_{x \leftarrow xs} T_f\ x) + |xs| + 1 \\
&T_{filter}\ T_p\ xs = (\textstyle\sum_{x \leftarrow xs} T_p\ x) + |xs| + 1 \\
&T_{take}\ n\ xs = min\ n\ |xs| + 1 \\
&T_{drop}\ n\ xs = min\ n\ |xs| + 1
\end{aligned}
$$

B.2 Selection

$T_{chop} :: nat \Rightarrow 'a\ list \Rightarrow nat$
$T_{chop}\ 0\ _ = 1$
$T_{chop}\ _\ [] = 1$
$T_{chop}\ n\ xs = T_{take}\ n\ xs + T_{drop}\ n\ xs + T_{chop}\ n\ (drop\ n\ xs) + 1$

$T_{partition3} :: 'a \Rightarrow 'a\ list \Rightarrow nat$
$T_{partition3}\ _\ [] = 1$
$T_{partition3}\ x\ (_ \#\ ys) = T_{partition3}\ x\ ys + 1$

$T_{slow_select} :: nat \Rightarrow 'a\ list \Rightarrow nat$
$T_{slow_select}\ k\ xs = T_{isort}\ xs + T_{nth}\ (isort\ xs)\ k + 1$

$T_{slow_median} :: 'a\ list \Rightarrow nat$
$T_{slow_median}\ xs = T_{slow_select}\ ((|xs| - 1)\ div\ 2)\ xs + 1$

$T_{chop}\ d\ xs \leq 5 \cdot |xs| + 1$
$T_{partition3}\ x\ xs = |xs| + 1$
$T_{slow_select}\ k\ xs \leq |xs|^2 + 3 \cdot |xs| + 3$
$T_{slow_median}\ xs \leq |xs|^2 + 3 \cdot |xs| + 4$

B.3 2-3 Trees

$T_{join_adj} :: 'a\ tree23s \Rightarrow nat$
$T_{join_adj}\ (TTs\ _\ _\ (T\ _)) = 1$
$T_{join_adj}\ (TTs\ _\ _\ (TTs\ _\ _\ (T\ _))) = 1$
$T_{join_adj}\ (TTs\ _\ _\ (TTs\ _\ _\ ts)) = T_{join_adj}\ ts + 1$

$T_{join_all} :: 'a\ tree23s \Rightarrow nat$
$T_{join_all}\ (T\ _) = 1$
$T_{join_all}\ ts = T_{join_adj}\ ts + T_{join_all}\ (join_adj\ ts) + 1$

$T_{tree23_of_list} :: 'a\ list \Rightarrow nat$
$T_{tree23_of_list}\ as = T_{leaves}\ as + T_{join_all}\ (leaves\ as) + 1$

B.4 Leftist Heaps

```
T_merge :: ('a × nat) tree ⇒ ('a × nat) tree ⇒ nat
T_merge ⟨⟩ t = 1
T_merge t ⟨⟩ = 1
T_merge (⟨l₁, (a₁, n₁), r₁⟩ =: t₁) (⟨l₂, (a₂, n₂), r₂⟩ =: t₂)
= (if a₁ ≤ a₂ then T_merge r₁ t₂ else T_merge t₁ r₂) + 1

T_insert :: 'a ⇒ ('a × nat) tree ⇒ nat
T_insert x t = T_merge ⟨⟨⟩, (x, 1), ⟨⟩⟩ t + 1

T_del_min :: ('a × nat) tree ⇒ nat
T_del_min ⟨⟩ = 1
T_del_min ⟨l, _, r⟩ = T_merge l r + 1
```

B.5 Binomial Heaps

```
T_link :: 'a tree ⇒ 'a tree ⇒ nat
T_link _ _ = 1

T_ins_tree :: 'a tree ⇒ 'a tree list ⇒ nat
T_ins_tree _ [] = 1
T_ins_tree t₁ (t₂ # ts)
= (if rank t₁ < rank t₂ then 1
   else T_link t₁ t₂ + T_ins_tree (link t₁ t₂) ts)

T_insert :: 'a ⇒ 'a tree list ⇒ nat
T_insert x ts = T_ins_tree (Node 0 x []) ts + 1

T_merge :: 'a tree list ⇒ 'a tree list ⇒ nat
T_merge _ [] = 1
T_merge [] (_ # _) = 1
T_merge (t₁ # ts₁ =: h₁) (t₂ # ts₂ =: h₂)
= 1 +
  (if rank t₁ < rank t₂ then T_merge ts₁ h₂
```

```
    else if rank t₂ < rank t₁ then T_merge h₁ ts₂
        else T_ins_tree (link t₁ t₂) (merge ts₁ ts₂) + T_merge ts₁ ts₂)

T_get_min :: 'a tree list ⇒ nat
T_get_min [_] = 1
T_get_min (_ # v # va) = 1 + T_get_min (v # va)

T_get_min_rest :: 'a tree list ⇒ nat
T_get_min_rest [_] = 1
T_get_min_rest (_ # v # va) = 1 + T_get_min_rest (v # va)

T_rev :: 'a list ⇒ nat
T_rev xs = |xs| + 1

T_del_min :: 'a tree list ⇒ nat
T_del_min ts
= T_get_min_rest ts +
  (case get_min_rest ts of
   (Node _ _ ts₁, ts₂) ⇒ T_rev ts₁ + T_merge (rev ts₁) ts₂) +
  1
```

B.6 Queues

```
T_norm :: 'a list × 'a list ⇒ nat
T_norm (fs, rs) = (if fs = [] then T_itrev rs [] else 0) + 1

T_enq :: 'a ⇒ 'a list × 'a list ⇒ nat
T_enq a (fs, rs) = T_norm (fs, a # rs) + 1

T_deq :: 'a list × 'a list ⇒ nat
T_deq (fs, rs) = (if fs = [] then 0 else T_norm (tl fs, rs)) + 1

T_first :: 'a list × 'a list ⇒ nat
T_first (_ # _, _) = 1
```

```
T_is_empty :: 'a list × 'a list ⇒ nat
T_is_empty (_, _) = 1
```

B.7 Splay Trees

```
T_splay :: 'a ⇒ 'a tree ⇒ nat
T_splay _ ⟨⟩ = 1
T_splay x ⟨AB, b, CD⟩
= (case cmp x b of
    LT ⇒ case AB of
           ⟨⟩ ⇒ 1 |
           ⟨A, a, B⟩ ⇒ case cmp x a of
                         LT ⇒ if A = ⟨⟩ then 1 else T_splay x A + 1 |
                         EQ ⇒ 1 |
                         GT ⇒ if B = ⟨⟩ then 1 else T_splay x B + 1 |
    EQ ⇒ 1 |
    GT ⇒ case CD of
           ⟨⟩ ⇒ 1 |
           ⟨C, c, D⟩ ⇒ case cmp x c of
                         LT ⇒ if C = ⟨⟩ then 1 else T_splay x C + 1 |
                         EQ ⇒ 1 |
                         GT ⇒ if D = ⟨⟩ then 1 else T_splay x D + 1)

T_splay_max :: 'a tree ⇒ nat
T_splay_max ⟨⟩ = 1
T_splay_max ⟨_, _, ⟨⟩⟩ = 1
T_splay_max ⟨_, _, ⟨_, _, C⟩⟩ = (if C = ⟨⟩ then 1 else T_splay_max C + 1)

T_insert :: 'a ⇒ 'a tree ⇒ nat
T_insert x t = 1 + (if t = ⟨⟩ then 0 else T_splay x t)

T_delete :: 'a ⇒ 'a tree ⇒ nat
T_delete x t
= 1 +
  (if t = ⟨⟩ then 0
```

```
    else T_splay x t +
         (case splay x t of
          ⟨l, a, _⟩ ⇒
            if x ≠ a then 0 else if l = ⟨⟩ then 0 else T_splay_max l))
```

B.8 Skew Heaps

```
T_merge :: 'a tree ⇒ 'a tree ⇒ nat
T_merge ⟨⟩ _ = 1
T_merge _ ⟨⟩ = 1
T_merge ⟨l₁, a₁, r₁⟩ ⟨l₂, a₂, r₂⟩
= (if a₁ ≤ a₂ then T_merge ⟨l₂, a₂, r₂⟩ r₁ else T_merge ⟨l₁, a₁, r₁⟩ r₂) + 1

T_insert :: 'a ⇒ 'a tree ⇒ int
T_insert a t = T_merge ⟨⟨⟩, a, ⟨⟩⟩ t + 1

T_del_min :: 'a tree ⇒ int
T_del_min t = (case t of ⟨⟩ ⇒ 1 | ⟨t₁, _, t₂⟩ ⇒ T_merge t₁ t₂ + 1)
```

B.9 Pairing Heaps

```
T_insert :: 'a ⇒ 'a tree ⇒ nat
T_insert _ _ = 1

T_merge :: 'a tree ⇒ 'a tree ⇒ nat
T_merge _ _ = 1

T_del_min :: 'a tree ⇒ nat
T_del_min ⟨⟩ = 1
T_del_min ⟨hs, _, _⟩ = T_pass2 (pass₁ hs) + T_pass1 hs + 1

T_pass1 :: 'a tree ⇒ nat
```

```
T_pass1 ⟨_, _, ⟨_, _, hs''⟩⟩ = T_pass1 hs'' + 1
T_pass1 ⟨⟩ = 1
T_pass1 ⟨_, _, ⟨⟩⟩ = 1

T_pass2 :: 'a tree ⇒ nat
T_pass2 ⟨⟩ = 1
T_pass2 ⟨_, _, hs⟩ = T_pass2 hs + 1
```

C

Notation

C.1 Symbol Table

The following table gives an overview of all the special symbols used in this book and how to enter them into Isabelle. The second column shows the full internal name of the symbol; the third column shows additional ASCII abbreviations. Either of these can be used to input the character using the auto-completion popup.

	Code	ASCII abbrev.	Comment
λ	`\<lambda>`	`%`	function abstraction
$\equiv$	`\<equiv>`	`==`	meta equality
$\neq$	`\<noteq>`	`~=`	
$\bigwedge$	`\<And>`	`!!`	meta $\forall$-quantifier
$\forall$	`\<forall>`	`!`	HOL $\forall$-quantifier
$\exists$	`\<exists>`	`?`	
$\Longrightarrow$	`\<Longrightarrow>`	`==>`	meta implication
$\longrightarrow$	`\<longrightarrow>`	`->`	HOL implication
$\longleftrightarrow$	`\<longleftrightarrow>`	`<->` or `<-->`	
$\Rightarrow$	`\<Rightarrow>`	`=>`	arrow in function types
$\leftarrow$	`\<leftarrow>`	`<-`	list comprehension syntax
$\neg$	`\<not>`	`~`	
$\wedge$	`\<and>`	`/\` or `&`	
$\vee$	`\<or>`	`\/` or `|`	
$\in$	`\<in>`	`:`	
$\notin$	`\<notin>`	`~:`	
$\bigcup$	`\<union>`	`Un`	

	Code	ASCII abbrev.	Comment
$\cap$	`\<inter>`	`Int`	
$\bigcup$	`\<Union>`	`Union or UN`	union/intersection of a set of sets
$\bigcap$	`\<Inter>`	`Inter or INT`	
$\subseteq$	`\<subseteq>`	`(=`	
$\subset$	`\<subset>`		
$\leq$	`\<le>`	`<=`	
$\geq$	`\<ge>`	`>=`	
$\circ$	`\<circ>`		function composition
$\times$	`\<times>`	`<*>`	cartesian prod., prod. type
$\vert$	`\<bar>`	`||`	absolute value
$\lfloor$	`\<lfloor>`	`[.`	floor
$\rfloor$	`\<rfloor>`	`.]`	
$\lceil$	`\<lceil>`	`[.`	ceiling
$\rceil$	`\<rceil>`	`.]`	
$\sum$	`\<Sum>`	`SUM`	see Section C.3
$\prod$	`\<Prod>`	`PROD`	

Note that the symbols "$\{\!\!\{$" and "$\}\!\!\}$" that is used for multiset notation in the book do not exist Isabelle; instead, the ASCII notation `{#` and `#}` are used (cf. Section C.3).

C.2 Subscripts and Superscripts

In addition to this, subscripts and superscripts with a single symbol can be rendered using two special symbols, `\<^sub>` and `\<^sup>`. The term x_0 for instance can be input as `x\<^sub>0`.

Longer subscripts and superscripts can be written using the symbols `\<^bsub>`...`\<^esub>` and `\<^bsup>`...`\<^esup>`, but this is only rendered in the somewhat visually displeasing form $\searrow \ldots \swarrow$ and $\nearrow \ldots \nwarrow$ by Isabelle/jEdit.

C.3 Syntactic Sugar

The following table lists relevant syntactic sugar that is used in the book or its supplementary material. In some cases, the book notation deviates slightly from the Isabelle notation for better readability.

The last column gives the formal meaning of the notation (i.e. what it expands to). In most cases, this is not important for the user to know, but it can occasionally be useful to find relevant lemmas, or to understand that e.g. if one encounters the term *sum f A*, this is just the η-contracted form of $\sum x \in A.\ f\ x$.

The variables in the table follow the following convention:

- x and y are of arbitrary type
- m and n are natural numbers
- P and Q are boolean values or predicates
- xs is a list
- A is a set
- M is a multiset

Book notation	Isabelle notation	Meaning	
	Arithmetic (for numeric types)		
$x \cdot y$	$x * y$	*times x y*	
$x\ /\ y$ or $\frac{x}{y}$	$x\ /\ y$	*divide x y*	(for type *real*)
x *div* y	x *div* y	*divide x y*	(for type *nat* or *int*)
$\lvert x \rvert$	$\lvert x \rvert$	*abs x*	
$\lfloor x \rfloor$	$\lfloor x \rfloor$	*floor x*	
$\lceil x \rceil$	$\lceil x \rceil$	*ceiling x*	
x^n	$x\ \hat{}\ n$	*power x n*	
	Lists		
$\lvert xs \rvert$		*length xs*	
$[]$	$[]$	*Nil*	
$x\ \#\ xs$	$x\ \#\ xs$	*Cons x xs*	
$[x, y]$	$[x, y]$	$x\ \#\ y\ \#\ []$	
$[m..<n]$	$[m..<n]$	*upt m n*	
$xs\ !\ n$	$xs\ !\ n$	*nth xs n*	
$xs[n := y]$	$xs[n := y]$	*list_update xs n y*	

Book notation	Isabelle notation	Meaning
	Sets	
$\{\}$	$\{\}$	*empty*
$\{x, y\}$	$\{x, y\}$	*insert x (insert y {})*
$x \in A$	$x \in A$	*Set.member x A*
$x \notin A$	$x \notin A$	$\neg(x \in A)$
$A \cup B$	$A \cup B$	*union A B*
$A \cap B$	$A \cap B$	*inter A B*
$A \subseteq B$	$A \subseteq B$	*subset_eq A B*
$A \subset B$	$A \subset B$	*subset A B*
$f \mathbin{`} A$	$f \mathbin{`} A$	*image f A*
$f -` A$	$f -` A$	*vimage f A*
$\{x \mid P\ x\}$	$\{x.\ P\ x\}$	*Collect P*
$\{x \in A \mid P\ x\}$	$\{x{\in}A.\ P\ x\}$	$\{x.\ P\ x \wedge x \in A\}$
$\{f\ x\ y \mid P\ x\ y\}$	$\{f\ x\ y \mid x\ y.\ P\ x\ y\}$	$\{z.\ \exists x\ y.\ z = f\ x\ y \wedge P\ x\ y\}$
$\bigcup_{x\in A} f\ x$	$\bigcup x{\in}A.\ f\ x$	$\bigcup(f \mathbin{`} A)$
$\forall x{\in}A.\ P\ x$	$\forall x{\in}A.\ P\ x$	*Ball A P*
$\exists x{\in}A.\ P\ x$	$\exists x{\in}A.\ P\ x$	*Bex A P*
	Multisets	
$\|M\|$		*size M*
$\{\!\{\}\!\}$	$\{\#\}$	*empty_mset*
$\{\!\{x\}\!\} + M$		*add_mset x M*
$\{\!\{x, y\}\!\}$	$\{\#x, y\#\}$	*add_mset x (add_mset y {#})*
$x \in_\# M$	$x \in\# M$	$x \in$ *set_mset M*
$x \notin_\# M$	$x \notin\# M$	$\neg(x \in\# M)$
$\{\!\{x \in_\# M \mid P\ x\}\!\}$	$\{\#\ x{\in}\#\ M.\ P\ x\ \#\}$	*filter_mset P M*
$\{\!\{f\ x \mid x \in_\# M\}\!\}$	$\{\#\ f\ x.\ x \in\# M\ \#\}$	*image_mset f M*
$\forall x{\in}_\# M.\ P\ x$	$\forall x{\in}\#M.\ P\ x$	$\forall x{\in}$*set_mset M*$.\ P\ x$
$\exists x{\in}_\# M.\ P\ x$	$\exists x{\in}\#M.\ P\ x$	$\exists x{\in}$*set_mset M*$.\ P\ x$
$M \subseteq_\# M'$	$M \subseteq\# M'$	*subseteq_mset M M'*

Book notation	Isabelle notation	Meaning
	Sums	
$\sum\ A$	$\sum A$	*sum* $(\lambda x.\ x)\ A$
$\sum_{x \in A}\ f\ x$	$\sum x{\in}A.\ f\ x$	*sum* $f\ A$
$\sum_{k\ =\ i}^{j}\ f\ k$	$\sum k{=}i..j.\ f\ k$	*sum* $f\ \{i..j\}$
$\sum_{\#}\ M$	$\sum_{\#} M$	*sum_mset* M
$\sum_{x \in_{\#} M}\ f\ x$	$\sum x{\in}\#M.\ f\ x$	*sum_mset* $(image_mset\ f\ M)$
$\sum_{x \leftarrow xs}\ f\ x$	$\sum x{\leftarrow}xs.\ f\ x$	*sum_list* $(map\ f\ xs)$
	(analogous for products)	
	Intervals (for ordered types)	
$\{x..\}$	$\{x..\}$	*atLeast* x
$\{..y\}$	$\{..y\}$	*atMost* y
$\{x..y\}$	$\{x..y\}$	*atLeastAtMost* $x\ y$
$\{x..{<}y\}$	$\{x..{<}y\}$	*atLeastLessThan* $x\ y$
$\{x{<}..y\}$	$\{x{<}..y\}$	*greaterThanAtMost* $x\ y$
$\{x{<}..{<}y\}$	$\{x{<}..{<}y\}$	*greaterThanLessThan* $x\ y$

References

1. S. Adams. Efficient sets - A balancing act. *J. Funct. Program.*, 3(4):553–561, 1993.
2. G. M. Adel'son-Vel'skiĭ and E. M. Landis. An algorithm for the organization of information. *Soviet Mathematics Doklady*, 3:1259–1263, 1962. Translated from Russian by M.J. Ricci.
3. M. Akra and L. Bazzi. On the solution of linear recurrence equations. *Computational Optimization and Applications*, 10(2):195–210, 1998.
4. A. Appel. Efficient verified red-black trees. Unpublished, 2011.
5. C. Ballarin. *Tutorial to Locales and Locale Interpretation.* https://isabelle.in.tum.de/doc/locales.pdf.
6. R. Bayer. Symmetric binary B-trees: Data structure and maintenance algorithms. *Acta Informatica*, 1:290–306, 1972.
7. S. Berghofer and M. Wenzel. Inductive datatypes in HOL - lessons learned in formal-logic engineering. In Y. Bertot, G. Dowek, A. Hirschowitz, C. Paulin-Mohring, and L. Théry, editors, *Theorem Proving in Higher Order Logics, TPHOLs'99*, volume 1690 of *LNCS*, pages 19–36. Springer, 1999.
8. J. C. Blanchette. Proof pearl: Mechanizing the textbook proof of Huffman's algorithm in Isabelle/HOL. *J. Autom. Reasoning*, 43(1):1–18, 2009.
9. G. E. Blelloch, D. Ferizovic, and Y. Sun. Just join for parallel ordered sets. In C. Scheideler and S. Gilbert, editors, *Proc. 28th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2016*, pages 253–264. ACM, 2016.
10. M. Blum, R. W. Floyd, V. Pratt, R. L. Rivest, and R. E. Tarjan. Time bounds for selection. *J. Comput. Syst. Sci*, 7(4):448–461, 1973.
11. F. W. Burton. An efficient functional implementation of FIFO queues. *Inf. Process. Lett.*, 14(5):205–206, 1982.
12. S. Cho and S. Sahni. Weight-biased leftist trees and modified skip lists. *ACM J. Exp. Algorithmics*, 3:2, 1998.
13. T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms, 3rd Edition.* MIT Press, 2009.
14. C. A. Crane. *Linear Lists and Priority Queues as Balanced Binary Trees.* PhD thesis, Stanford University, 1972. STAN-CS-72-259.

15. R. De La Briandais. File searching using variable length keys. In *Papers Presented at the the March 3-5, 1959, Western Joint Computer Conference*, IRE-AIEE-ACM '59 (Western), pages 295–298. ACM, 1959.
16. M. Eberl. The number of comparisons in quicksort. *Archive of Formal Proofs*, Mar. 2017. http://isa-afp.org/entries/Quick_Sort_Cost.html, Formal proof development.
17. M. Eberl. Proving divide and conquer complexities in Isabelle/HOL. *Journal of Automated Reasoning*, 58(4):483–508, 2017.
18. M. Eberl, M. W. Haslbeck, and T. Nipkow. Verified analysis of random binary tree structures. In J. Avigad and A. Mahboubi, editors, *Interactive Theorem Proving (ITP 2018)*, volume 10895 of *LNCS*, pages 196–214. Springer, 2018.
19. J.-C. Filliâtre and P. Letouzey. Functors for proofs and programs. In *ESOP*, volume 2986 of *LNCS*, pages 370–384. Springer, 2004.
20. E. Fredkin. Trie memory. *Commun. ACM*, 3(9):490–499, 1960.
21. M. L. Fredman, R. Sedgewick, D. Sleator, and R. Tarjan. The pairing heap: A new form of self-adjusting heap. *Algorithmica*, 1(1):111–129, 1986.
22. K. Germane and M. Might. Deletion: The curse of the red-black tree. *J. Functional Programming*, 24(4):423–433, 2014.
23. L. J. Guibas and R. Sedgewick. A dichromatic framework for balanced trees. In *19th Annual Symposium on Foundations of Computer Science (FOCS 1978)*, pages 8–21, 1978.
24. F. Haftmann. *Code generation from Isabelle/HOL theories*. http://isabelle.in.tum.de/doc/codegen.pdf.
25. F. Haftmann. *Haskell-style type classes with Isabelle/Isar*. http://isabelle.in.tum.de/doc/classes.pdf.
26. F. Haftmann and T. Nipkow. Code generation via higher-order rewrite systems. In M. Blume, N. Kobayashi, and G. Vidal, editors, *Functional and Logic Programming (FLOPS 2010)*, volume 6009 of *LNCS*, pages 103–117. Springer, 2010.
27. Haskell website. https://www.haskell.org.
28. R. Hinze. On constructing 2-3 trees. *J. Funct. Program.*, 28:e19, 2018.
29. C. M. Hoffmann and M. J. O'Donnell. Programming with equations. *ACM Trans. Program. Lang. Syst.*, 4(1):83–112, 1982.
30. R. Hood and R. Melville. Real-time queue operation in pure LISP. *Inf. Process. Lett.*, 13(2):50–54, 1981.
31. R. R. Hoogerwoord. A logarithmic implementation of flexible arrays. In R. Bird, C. Morgan, and J. Woodcock, editors, *Mathematics of Program Construction*, volume 669 of *LNCS*, pages 191–207. Springer, 1992.
32. D. A. Huffman. A method for the construction of minimum-redundancy codes. In *Proceedings of the I.R.E.*, pages 1098–1101, 1952.
33. J. Iacono. Improved upper bounds for pairing heaps. In M. M. Halldórsson, editor, *Algorithm Theory - SWAT 2000, 7th Scandinavian Workshop on Algorithm Theory*, volume 1851 of *LNCS*, pages 32–45. Springer, 2000.
34. J. Iacono and M. V. Yagnatinsky. A linear potential function for pairing heaps. In T. H. Chan, M. Li, and L. Wang, editors, *Combinatorial Optimization and Applications — 10th International Conference, COCOA 2016*, volume 10043 of *LNCS*, pages 489–504. Springer, 2016.

35. C. B. Jones. *Systematic Software Development using VDM*. Prentice Hall International, 2nd edition, 1990.
36. S. Kahrs. Red black trees with types. *J. Functional Programming*, 11(4):425–432, 2001.
37. A. Kaldewaij and B. Schoenmakers. The derivation of a tighter bound for top-down skew heaps. *Information Processing Letters*, 37:265–271, 1991.
38. ACM Paris Kanellakis Theory and Practice Award. https://awards.acm.org/kanellakis.
39. R. M. Karp. Probabilistic recurrence relations. *J. ACM*, 41(6):11361150, Nov. 1994.
40. D. J. King. Functional binomial queues. In K. Hammond, D. N. Turner, and P. M. Sansom, editors, *Proceedings of the 1994 Glasgow Workshop on Functional Programming*, Workshops in Computing, pages 141–150. Springer, 1994.
41. D. E. Knuth. Optimum binary search trees. *Acta Informatica*, 1:14–25, 1971.
42. D. E. Knuth. Huffman's algorithm via algebra. *J. Comb. Theory, Ser. A*, 32(2):216–224, 1982.
43. D. E. Knuth. *The Art of Computer Programming, vol. 1: Fundamental Algorithms*. Addison–Wesley, 3rd edition, 1997.
44. A. Krauss. *Defining Recursive Functions in Isabelle/HOL*. http://isabelle.in.tum.de/doc/functions.pdf.
45. A. Krauss. Partial recursive functions in higher-order logic. In U. Furbach and N. Shankar, editors, *Automated Reasoning,IJCAR 2006*, volume 4130 of *LNCS*, pages 589–603. Springer, 2006.
46. P. Lammich and T. Nipkow. Proof Pearl: Purely Functional, Simple and Efficient Priority Search Trees and Applications to Prim and Dijkstra. In J. Harrison, J. O'Leary, and A. Tolmach, editors, *10th International Conference on Interactive Theorem Proving (ITP 2019)*, volume 141 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 23:1–23:18. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2019.
47. D. H. Larkin, S. Sen, and R. E. Tarjan. A back-to-basics empirical study of priority queues. In C. C. McGeoch and U. Meyer, editors, *2014 Proceedings of the Sixteenth Workshop on Algorithm Engineering and Experiments, ALENEX 2014*, pages 61–72. SIAM, 2014.
48. T. Leighton. Notes on better master theorems for divide-and-conquer recurrences (MIT lecture notes). 1996.
49. R. Meis, F. Nielsen, and P. Lammich. Binomial heaps and skew binomial heaps. *Archive of Formal Proofs*, Oct. 2010. http://isa-afp.org/entries/Binomial-Heaps.html, Formal proof development.
50. D. R. Morrison. PATRICIA - practical algorithm to retrieve information coded in alphanumeric. *J. ACM*, 15(4):514–534, 1968.
51. P. Müller. The binomial heap verification challenge in Viper. In P. Müller and I. Schaefer, editors, *Principled Software Development*, pages 203–219. Springer, 2018.
52. T. Nipkow. *Programming and Proving in Isabelle/HOL*. http://isabelle.in.tum.de/doc/prog-prove.pdf.
53. T. Nipkow. Amortized complexity verified. In C. Urban and X. Zhang, editors, *Interactive Theorem Proving (ITP 2015)*, volume 9236 of *LNCS*, pages 310–324. Springer, 2015.

54. T. Nipkow. Automatic functional correctness proofs for functional search trees. In J. Blanchette and S. Merz, editors, *Interactive Theorem Proving (ITP 2016)*, volume 9807 of *LNCS*, pages 307–322. Springer, 2016.
55. T. Nipkow and H. Brinkop. Amortized complexity verified. *J. Autom. Reasoning*, 62(3):367–391, 2019.
56. T. Nipkow and G. Klein. *Concrete Semantics with Isabelle/HOL*. Springer, 2014. 298 pp. http://concrete-semantics.org.
57. T. Nipkow, L. Paulson, and M. Wenzel. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer, 2002. http://www.in.tum.de/~nipkow/LNCS2283/.
58. T. Nipkow and T. Sewell. Proof pearl: Braun trees. In J. Blanchette and C. Hritcu, editors, *Certified Programs and Proofs, CPP 2020*, pages 18–31. ACM, 2020.
59. T. Nipkow and D. Somogyi. Optimal binary search trees. *Archive of Formal Proofs*, May 2018. https://isa-afp.org/entries/Optimal_BST.html, Formal proof development.
60. Ocaml website. https://ocaml.org.
61. C. Okasaki. Three algorithms on Braun trees. *J. Functional Programming*, 7(6):661–666, 1997.
62. C. Okasaki. *Purely Functional Data Structures*. Cambridge University Press, 1998.
63. L. C. Paulson. The foundation of a generic theorem prover. *Journal of Automated Reasoning*, 5:363–397, 1989.
64. L. C. Paulson. *ML for the Working Programmer*. Cambridge University Press, 2nd edition, 1996.
65. S. Pettie. Towards a final analysis of pairing heaps. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS*, pages 174–183. IEEE Computer Society, 2005.
66. C. Reade. Balanced trees with removals: An exercise in rewriting and proof. *Sci. Comput. Program.*, 18(2):181–204, 1992.
67. M. Rem and W. Braun. A logarithmic implementation of flexible arrays. Memorandum MR83/4. Eindhoven University of Techology, 1983.
68. S. Sahni. Leftist trees. In D. P. Mehta and S. Sahni, editors, *Handbook of Data Structures and Applications*. Chapman and Hall/CRC, 2004.
69. B. Schoenmakers. A systematic analysis of splaying. *Information Processing Letters*, 45:41–50, 1993.
70. D. D. Sleator and R. E. Tarjan. Self-adjusting binary search trees. *J. ACM*, 32(3):652–686, 1985.
71. D. D. Sleator and R. E. Tarjan. Self-adjusting heaps. *SIAM J. Comput.*, 15(1):52–69, 1986.
72. R. E. Tarjan. Amortized complexity. *SIAM J. Alg. Disc. Meth.*, 6(2):306–318, 1985.
73. L. Théry. Formalising Huffman's algorithm. Technical Report TRCS 034, Department of Informatics, University of L'Aquila, 2004.
74. F. Turbak. CS230 Handouts — Spring 2007, 2007. http://cs.wellesley.edu/~cs230/spring07/handouts.html.
75. J. Vuillemin. A data structure for manipulating priority queues. *Commun. ACM*, 21(4):309–315, 1978.

76. M. Wenzel. *Isabelle/Isar — A Versatile Environment for Human-Readable Formal Proof Documents*. PhD thesis, Institut für Informatik, Technische Universität München, 2002. http://nbn-resolving.de/urn/resolver.pl?urn:nbn:de:bvb:91-diss2002020117092.
77. J. Williams. Algorithm 232 — Heapsort. *Communications of the ACM*, 7(6):347–348, 1964.
78. S. Wimmer, S. Hu, and T. Nipkow. Monadification, memoization and dynamic programming. *Archive of Formal Proofs*, May 2018. https://isa-afp.org/entries/Monad_Memo_DP.html, Formal proof development.
79. S. Wimmer, S. Hu, and T. Nipkow. Verified memoization and dynamic programming. In J. Avigad and A. Mahboubi, editors, *Interactive Theorem Proving (ITP 2018)*, volume 10895 of *Lecture Notes in Computer Science*, pages 579–596. Springer, 2018.
80. F. F. Yao. Efficient dynamic programming using quadrangle inequalities. In *STOC*, pages 429–435. ACM, 1980.
81. B. Zhan. Efficient verification of imperative programs using auto2. In D. Beyer and M. Huisman, editors, *Tools and Algorithms for the Construction and Analysis of Systems, TACAS 2018*, volume 10805 of *LNCS*, pages 23–40. Springer, 2018.

Index

|_|$_1$, 50
|_|, 50, 251
[], 8
[_ ..< _], 252
(@), 251
(!), 252
(#), 8
_ [_ := _], 252
{#}, 10
$\sum_{\#}$, 10
($\in_{\#}$), 10
⟨_, _, _⟩, 49
⟨⟩, 49
2-3 trees, 81

abstraction function, 61, 76
acomplete, 53
almost complete tree, 53
amortized running time, 214
array, 121
augmented tree, 57
AVL tree, 101

balance factor, 108
balanced search trees, 81
Bellman-Ford algorithm, 202
binary search trees, 61
binary trie, 145
binomial heap, 179
black height, 92
Braun tree, 122, 123
bst, 61
butlast, 252

child, 49
cmp, 62
coercion, 9
complete, 52
complete tree, 51
computation induction, 12
concat, 251

deletion by joining, 64
deletion by replacing, 64
depth, 49
difference lists, 30
drop, 252
dynamic programming, 191

Fibonacci tree, 107
filter, 251
filter_mset, 10

h, 51
hd, 252
head, 8
heap, 168
height, 51
height-balanced tree, 108
Huffman's algorithm, 151

IH, 12
image_mset, 10
inclusion, 9
inorder, 50

insertion sort, 20
int, 9
interval tree, 69
invariant, 65, 76
  preservation, 65

join approach, 114

leftist heap, 171
leftist tree, 171
length, 251
level, 49
linear order, 19
linorder, 19
list, 8
list extensionality, 127
locales, 76

Map, 78
map, 251
map_tree, 50
median, 40
merge sort, 24
mh, 51
min_height, 51
min_list, 253
mset, 10
multiset, 10
multiset, 10

nat, 9

observer functions, 215
optimal binary search tree, 207

pairing heap, 241
Patricia trie, 147
pivot, 22
preorder, 50
Priority_Queue, 167
priority queue, 167
  mergeable, 167
Priority_Queue_Merge, 168

Queue, 217
queue, 217
  Hood-Melville, 219
quicksort, 22

real, 9
red-black tree, 91
replicate, 252
rev, 252
right-heavy, 238
root, 49
running time, 12
running time function, 12
  canonical, 14
runs, 29

selection, 35
Set, 76
set, 251
set_mset, 10
set_tree, 50
skew heap, 237
sorted, 19, 66
sorted_wrt, 253
sorting, 19
specification, 65
splay tree, 229
stability (sorting), 33
stable, 33
sum_list, 253

T, 12
tail, 8
take, 252
time function, *see* running time function
tl, 252
tree, 49
trie, 143
2-3 trees, 81
type class, 19
type of interest, 76
type variable, 8

unit, 9