

Automatic Functional Correctness Proofs for Functional Search Trees

Tobias Nipkow^(✉)

Technische Universität München, Munich, Germany
nipkow@in.tum.de
<http://www.in.tum.de/~nipkow>

Abstract. In a new approach, functional correctness specifications of *insert/update* and *delete* operations on search trees are expressed on the level of lists by means of an inorder traversal function that projects trees to lists. With the help of a small lemma library, functional correctness and preservation of the search tree property are proved automatically (in Isabelle/HOL) for a range of data structures: unbalanced binary trees, AVL trees, red-black trees, 2-3 and 2-3-4 trees, 1-2 brother trees, AA trees and splay trees.

1 Introduction

Most books and articles on search tree data structures do not discuss functional correctness, which is taken to be obvious, but concentrate on non-obvious structural invariants like balancedness. This paper confirms that this is the right attitude by providing a framework for proving the functional correctness of eight different search tree data structures automatically (in Isabelle/HOL [19,21]).

What is proved automatically? Functional correctness of *insert*, *delete* and *isin* together with the preservation of the search tree invariant, i.e. sortedness, by *insert* and *delete*. Structural invariants like balancedness are proved manually, depend on the specific data structure, and are not discussed here.

Which data structures are covered? Unbalanced binary trees, AVL trees, red-black trees, 2-3 and 2-3-4 trees, 1-2 brother trees, AA trees and splay trees.¹ As far as we know, these are the first formal proofs for 2-3 and 2-3-4 trees, 1-2 brother trees and AA trees, and the first automatic proofs for most of the eight data structures.

What does automatic mean? It means that all the required theorems are proved by induction followed by a single invocation of Isabelle's `auto` proof method, parameterized with a fixed set of basic lemmas plus further lemmas about auxiliary functions. The lemmas to be proved about *insert*, *delete* and *isin* are fixed; lemmas about auxiliary functions need to be invented but (mostly) follow a simple pattern.

T. Nipkow—Supported by DFG Koselleck grant NI 491/16-1.

¹ See http://isabelle.in.tum.de/library/HOL/HOL-Data_Structures/ or the source directory `src/HOL/Data_Structures/` in the Isabelle distribution.

The paper is structured as follows. Section 3 presents two approaches to the specification and verification of set implementations: the standard approach and our new approach. Section 4 details the verification framework behind the new approach. In Sect. 5 eight different search tree implementations and their correctness proofs are discussed. In a final section it is shown how the framework can be generalized from sets to maps.

Related work is discussed in the body of the paper. With one exception, the proofs in previous work are not automatic. We refrain from stating this each time and we do not describe how far from automatic they are, although this varies significantly (from a few to more than a hundred lines).

2 Lists and Trees

Lists (type $'a\ list$) are constructed from the empty list $[]$ via the infix constructor “.”. The notation $[x,y,z]$ is short for $x \cdot y \cdot z \cdot []$. The infix $@$ concatenates two lists.

Binary trees are defined as the data type $'a\ tree$ with two constructors: the empty tree or leaf $\langle \rangle$ and the node $\langle l, a, r \rangle$ with subtrees $l, r :: 'a\ tree$ and contents $a :: 'a$.

There is also a type $'a\ set$ of sets with their usual operations.

3 Set Implementations

We require that an implementation of sets provides some type $'a\ t$ (where $'a$ is the element type) and the operations

$$\begin{aligned} empty &:: 'a\ t \\ insert &:: 'a \Rightarrow 'a\ t \Rightarrow 'a\ t \\ delete &:: 'a \Rightarrow 'a\ t \Rightarrow 'a\ t \\ isin &:: 'a\ t \Rightarrow 'a \Rightarrow bool \end{aligned}$$

In the rest of the paper we ignore *empty* because it is trivial.

In order to specify these operations we assume that there is also an *abstraction function* $set :: 'a\ t \Rightarrow 'a\ set$ and a data type invariant $invar :: 'a\ t \Rightarrow bool$. These are not part of the interface and need not be executable but have to be provided in order to prove an implementation correct w.r.t. the specification in Fig. 1. Specifications phrased in terms of abstraction functions that are required to be homomorphisms go back to Hoare [11] and became an integral part of the model-oriented specification language VDM [13]. In the first-order context of universal algebra it was shown that there are always fully abstract models such that any concrete implementation can be shown correct with a homomorphism [16]. In Isabelle, implementations that satisfy such specifications can automatically be plugged in for the abstract type by regarding the abstraction function as a constructor [9]. For example, turning the equation $isin\ s\ x = (x \in set\ s)$ around tells us how to evaluate $x \in set\ s$ with the help of *isin*.

From now on we assume that the element type 'a is linearly ordered.

$$\begin{aligned}
\text{invar } s &\implies \text{set } (\text{insert } x \ s) = \{x\} \cup \text{set } s \\
\text{invar } s &\implies \text{set } (\text{delete } x \ s) = \text{set } s - \{x\} \\
\text{invar } s &\implies \text{isin } s \ x = (x \in \text{set } s) \\
\text{invar } s &\implies \text{invar } (\text{insert } x \ s) \\
\text{invar } s &\implies \text{invar } (\text{delete } x \ s)
\end{aligned}$$

Fig. 1. Specification of set implementations

3.1 The Standard Approach

The most compact form of the standard approach to the verification of search tree implementations of sets consists of the following items:

- An abstraction function *set* that extract the set of elements in a tree.
- A recursively defined (binary) search tree invariant

$$\text{bst } \langle l, a, r \rangle = (\text{bst } l \wedge \text{bst } r \wedge (\forall x \in \text{set } l. x < a) \wedge (\forall x \in \text{set } r. a < x))$$
- Proof of the correctness conditions in Fig. 1 where *invar* is *bst*, possibly conjoined with additional structural invariants.

There are many variations of the above setup, some of which address two complications that arise when automating the proofs, the quantifiers and the non-free data type of sets:

- In the definition of *bst*, quantifiers are replaced by auxiliary functions that check if all elements in a tree are less/greater than a given element.
- Instead of extensional equality of sets, e.g. $\text{set } (\text{insert } x \ s) = \{x\} \cup \text{set } s$, pointwise equality is proved, e.g. $\text{isin } (\text{insert } x \ s) \ y = (x = y \vee \text{isin } s \ y)$.
- The predicate *bst* is defined inductively rather than recursively.

We subsume all of these variations under the “standard approach”. Unless stated otherwise, all related work follows the standard approach.

3.2 The *inorder* Approach

Why is it perfectly obvious that the following equation (where *R/B* construct red/black nodes) preserves sortedness and the set of elements of a tree?

$$\text{balance } (R (R \ t_1 \ a \ t_2) \ b \ t_3) \ c \ t_4 = R (B \ t_1 \ a \ t_2) \ b (B \ t_3 \ c \ t_4)$$

Because the sequence of subtrees and elements is the same on both sides! We merely need to make the machine see this as well as we do.

The key idea of our approach is to base it on the *inorder* traversal of trees. That is, we use lists as an intermediate data type between sets and trees. To this end we need four auxiliary functions on lists:

- $\uparrow :: 'a \ \text{list} \Rightarrow \text{bool}$
- \uparrow means that the list is sorted in ascending order w.r.t. $<$.

$$\begin{aligned}
\text{invar } t &\Longrightarrow \lfloor \text{insert } x \ t \rfloor = \text{ins}_{\text{list}} \ x \ \lfloor t \rfloor & (1) \\
\text{invar } t &\Longrightarrow \lfloor \text{delete } x \ t \rfloor = \text{del}_{\text{list}} \ x \ \lfloor t \rfloor & (2) \\
\text{invar } t &\Longrightarrow \text{isin } t \ x = (x \in \text{elems } \lfloor t \rfloor) & (3) \\
\text{invar } t &\Longrightarrow \text{inv } (\text{insert } x \ t) & (4) \\
\text{invar } t &\Longrightarrow \text{inv } (\text{delete } x \ t) & (5)
\end{aligned}$$

Fig. 2. Specification of set implementations over ordered types

$$\begin{aligned}
&\uparrow [] = \text{True} \\
&\uparrow [x] = \text{True} \\
&\uparrow (x \cdot y \cdot xs) = (x < y \wedge \uparrow (y \cdot xs)) \\
- \text{ins}_{\text{list}} &:: 'a \Rightarrow 'a \text{ list} \Rightarrow 'a \text{ list} \\
&\text{ins}_{\text{list}} \text{ inserts an element at the correct position into a sorted list if the element} \\
&\text{is not present in the list yet.} \\
&\text{ins}_{\text{list}} \ x \ [] = [x] \\
&\text{ins}_{\text{list}} \ x \ (a \cdot xs) = \\
&\quad (\text{if } x < a \ \text{then } x \cdot a \cdot xs \ \text{else if } x = a \ \text{then } a \cdot xs \ \text{else } a \cdot \text{ins}_{\text{list}} \ x \ xs) \\
- \text{del}_{\text{list}} &:: 'a \Rightarrow 'a \text{ list} \Rightarrow 'a \text{ list} \\
&\text{del}_{\text{list}} \text{ deletes the first occurrence of an element from a list.} \\
&\text{del}_{\text{list}} \ x \ [] = [] \\
&\text{del}_{\text{list}} \ x \ (a \cdot xs) = (\text{if } x = a \ \text{then } xs \ \text{else } a \cdot \text{del}_{\text{list}} \ x \ xs) \\
- \text{elems} &:: 'a \text{ list} \Rightarrow 'a \text{ set} \\
&\text{elems} \text{ turns a list into the set of its elements.} \\
&\text{elems } [] = \emptyset \\
&\text{elems } (x \cdot xs) = \{x\} \cup \text{elems } xs
\end{aligned}$$

A new specification of sets (over a linearly ordered type $'a$) is shown in Fig. 2. The crucial ingredient is an additional specification function

$$\text{inorder} :: 'a \ t \Rightarrow 'a \ \text{list}$$

As the name *inorder* suggests, you should now think of $'a \ t$ as a type of trees. We abbreviate $\text{inorder } t$ by $\lfloor t \rfloor$.

The fact that some t is a search tree, i.e. sorted, can now be expressed as $\uparrow \lfloor t \rfloor$. This invariant will be dealt with automatically. Of course search trees frequently have additional structural invariants. These can be supplied via yet another specification function $\text{inv} :: 'a \ t \Rightarrow \text{bool}$. Both kinds of invariants are combined into *invar*:

$$\text{invar } t = (\text{inv } t \wedge \uparrow \lfloor t \rfloor)$$

The first three propositions in Fig. 2 demand the functional correctness of *insert*, *delete* and *isin* w.r.t. ins_{list} , del_{list} and *elems*. The next two propositions demand that *inv* is invariant. If we interpret function *set* as $\text{elems} \circ \text{inorder}$ it is easy to show that Fig. 2 implies Fig. 1 with the help of the following simple inductive lemmas:

$$\begin{aligned}
elems (ins_{list} x xs) &= \{x\} \cup elems xs \\
\uparrow xs &\Longrightarrow distinct xs \\
distinct xs &\Longrightarrow elems (del_{list} x xs) = elems xs - \{x\} \\
\uparrow xs &\Longrightarrow \uparrow (ins_{list} x xs) \\
\uparrow xs &\Longrightarrow \uparrow (del_{list} x xs)
\end{aligned}$$

In summary: the functional correctness of an implementation of *insert*, *delete* and *isin* on some data structure can be verified by proving the properties in Fig. 2 for some suitable definition of *inorder* and *inv*. In the following section we introduce a library of lemmas about \uparrow , ins_{list} and del_{list} that allows us to automate the proofs of (1)–(3).

Note that although we equate (1)–(3) with “functional correctness”, it is more: (1)–(3) also imply that sortedness is an invariant.

4 The Verification Framework

We do not claim to provide a framework that can prove any implementation of sets by search trees automatically correct. Instead we provide lemmas that work in practice (they automate the correctness proofs for a list of benchmark implementations presented in Sect. 5) and are well motivated by general considerations concerning the shape of formulas that arise in the verification.

As a motivating example we consider ordinary unbalanced binary trees *'a tree*. The textbook definitions of *insert*, *delete* and *isin* are omitted. Let us examine how to prove

$$\uparrow [t] \Longrightarrow [insert\ x\ t] = ins_{list}\ x\ [t]$$

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:

$$\begin{aligned}
[insert\ x\ t] &= [insert\ x\ l] @ a \cdot [r] = ins_{list}\ x\ [l] @ a \cdot [r] \\
&= ins_{list}\ x\ ([l] @ a \cdot [r]) = ins_{list}\ x\ t
\end{aligned}$$

The first and last step are by definition, the second step by induction hypothesis, but the third step requires two lemmas:

$$\begin{aligned}
\uparrow (xs @ y \cdot ys) &= (\uparrow (xs @ [y]) \wedge \uparrow (y \cdot ys)) \\
\uparrow (xs @ [a]) \wedge x < a &\Longrightarrow ins_{list}\ x\ (xs @ a \cdot ys) = ins_{list}\ x\ xs @ a \cdot ys
\end{aligned}$$

The first lemma rewrites the assumption $\uparrow [t]$ to $\uparrow ([l] @ [a]) \wedge \uparrow (a \cdot [r])$, thus allowing the second lemma to rewrite the term $ins_{list}\ x\ ([l] @ a \cdot [r])$ to $ins_{list}\ x\ [l] @ a \cdot [r]$.

It may seem that the two lemmas just shown are rather arbitrary, but we will see that in the context of trees, where each node is a tuple $\langle s_0, a_1, s_1, \dots, s_n \rangle$ of subtrees s_i alternating with elements a_i , there is an underlying principle. In the properties in Fig. 2 the following three terms are crucial: $\uparrow [t]$, $ins_{list}\ x\ [t]$ and $del_{list}\ x\ [t]$. Assuming that the properties are proved by induction, t will be some (possibly complicated) tree constructor term. Evaluating $[t]$ will thus lead to a list of the following form where sublists and individual elements alternate:

$$\begin{aligned}
\uparrow (xs @ y \cdot ys) &= (\uparrow (xs @ [y]) \wedge \uparrow (y \cdot ys)) & (6) \\
\uparrow (x \cdot xs @ y \cdot ys) &= (\uparrow (x \cdot xs) \wedge x < y \wedge \uparrow (xs @ [y]) \wedge \uparrow (y \cdot ys)) & (7) \\
\uparrow (x \cdot xs) &\Longrightarrow \uparrow xs & (8) \\
\uparrow (xs @ [y]) &\Longrightarrow \uparrow xs & (9) \\
\uparrow (xs @ [a]) &\Longrightarrow ins_{list} x (xs @ a \cdot ys) = & (10) \\
&\quad (if\ x < a\ then\ ins_{list}\ x\ xs\ @\ a \cdot ys\ else\ xs\ @\ ins_{list}\ x\ (a \cdot ys)) \\
\uparrow (xs @ a \cdot ys) &\Longrightarrow del_{list} x (xs @ a \cdot ys) = & (11) \\
&\quad (if\ x < a\ then\ del_{list}\ x\ xs\ @\ a \cdot ys\ else\ xs\ @\ del_{list}\ x\ (a \cdot ys)) \\
elems (xs @ ys) &= elems xs \cup elems ys & (12) \\
\uparrow (y \cdot xs) \wedge x \leq y &\Longrightarrow x \notin elems xs & (13) \\
\uparrow (xs @ [y]) \wedge y \leq x &\Longrightarrow x \notin elems xs & (14)
\end{aligned}$$

Fig. 3. Lemmas for \uparrow , ins_{list} , del_{list} and $elems$

$$[t_1] @ a_1 \cdot [t_2] @ a_2 \cdot \dots \cdot [t_n]$$

Now we discuss a set of lemmas (see Fig. 3) that allow us to simplify the application of \uparrow , ins_{list} and del_{list} to such terms.

Terms of the form $\uparrow(xs_1 @ a_1 \cdot xs_2 @ a_2 \cdot \dots \cdot xs_n)$ are decomposed into the following *basic* formulas

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

by the rewrite rules (6)–(7). Lemmas (8)–(9) enable deductions from basic formulas.

Terms of the form $ins_{list} x (xs_1 @ a_1 \cdot xs_2 @ a_2 \cdot \dots \cdot xs_n)$ are rewritten with equation (10) (and the defining equations for ins_{list}) to push ins_{list} inwards. Terms of the form $del_{list} x (xs_1 @ a_1 \cdot xs_2 @ a_2 \cdot \dots \cdot xs_n)$ are rewritten with Eq. (11) (and the defining equations for del_{list}) to push del_{list} inwards.

Finally we need lemmas (12)–(14) about *elems* on sorted lists.

The lemmas in Fig. 3 form the complete set of basic lemmas on which the automatic proofs of almost all search trees in the paper rest; only splay trees need additional lemmas.

4.1 Proof Automation by Rewriting

The automatic proofs rely on conditional, contextual term rewriting with the following bells and whistles (which Isabelle’s simplifier provides):

- Conjunctions in the context are split up into their conjuncts.
- Conditionals and case-expressions can be split automatically.
- A decision procedure for linear orders that can decide if some literal (a possibly negated atom $a < b$ or $a \leq b$) follows from a set of literals in the context.

- Implications (8)–(9) lead to nontermination when used as conditional rewrite rules. It must be possible to direct the simplifier to solve the preconditions of those rules by assumptions in the context rather than a recursive simplifier invocation. In Isabelle there is a constant $ASSUMPTION = (\lambda x. x)$ that can be wrapped around a precondition of a rewrite rule and prevents recursive applications of the simplifier to that precondition.

5 An Arboretum

In the rest of this section we focus on (1) and (2) when discussing the proofs of the properties in Fig. 2. This is because requirement (3) can always (except for splay trees) be proved automatically without further lemmas and (4) and (5) are specific to the individual data structures and not part of functional correctness.

Because there is not enough space to present all definitions and proofs, Table 1 gives an overview in terms of lines of code and numbers of functions needed for each data structure. Because *isin* is (almost) the same for all of them (except splay trees), it is excluded. The table shows that there is at most one lemma per function, except for splay trees.

Table 1. Code and proof statistics for *insert + delete* (l.o. = lines of)

	Unbal	AVL	Red-Black	2-3	2-3-4	Brother	AA	Splay
l.o. code	17	45	61	88	143	66	55	46
functions	3	8	11	12	16	10	8	4
lemmas	3	6	11	10	14	10	6	5

The majority of lemmas about auxiliary functions follow a simple pattern. Typical examples are balancing functions, e.g. $[bal\ t] = [t]$, or smart constructors, e.g. $[node\ l\ a\ r] = [l] @ a \cdot [r]$. We call these *trivial lemmas*. More complicated lemmas are discussed explicitly in the text; we call them *non-trivial*.

All our implementations compare elements with a comparison operator *cmp* that returns an element of the **datatype** $cmp = LT \mid EQ \mid GT$.

5.1 Unbalanced Trees

Function *insert* is trivial and (1) is proved directly. Function *delete* is more interesting because it is defined with the help of an auxiliary function:

$$\begin{aligned}
 & delete\ x\ \langle \rangle = \langle \rangle \\
 & delete\ x\ \langle l, a, r \rangle = \\
 & (case\ cmp\ x\ a\ of\ LT \Rightarrow \langle delete\ x\ l, a, r \rangle \\
 & \mid EQ \Rightarrow if\ r = \langle \rangle\ then\ l\ else\ let\ (x, y) = del_min\ r\ in\ \langle l, x, y \rangle \\
 & \mid GT \Rightarrow \langle l, a, delete\ x\ r \rangle)
 \end{aligned}$$

$$\begin{aligned} del_min \langle l, a, r \rangle = \\ (if\ l = \langle \rangle\ then\ (a, r)\ else\ let\ (x, l') = del_min\ l\ in\ (x, \langle l', a, r \rangle)) \end{aligned}$$

The proof of (2) requires the following lemma about *del_min* that the user has to formulate himself; the proof is again automatic.

$$del_min\ t = (x, t') \wedge t \neq \langle \rangle \implies x \cdot [t'] = [t]$$

This is one of the more “difficult” lemmas to invent.

5.2 AVL Trees

Our starting point was an existing formalization [20] which follows the standard approach. Functional correctness of AVL trees can be proved without assuming any structural (height) invariants. The only non-trivial lemma we require is

$$del_max\ t = (t', a) \wedge t \neq \langle \rangle \implies [t'] @ [a] = [t]$$

Related Work. Filliâtre and Letouzey [6] report on a verification of AVL trees in Coq. They follow the standard approach, except that the executable functions are extracted from constructive proofs. An updated version of their proofs in the Coq distribution gives the functions explicitly. Ralston [26] reports a proof with ACL2. The verification by Clochard [4] in Why3 is interesting because he also abstracts trees to their inorder traversal and reports that the proofs for AVL trees are automatic.

5.3 Red-Black Trees

Red-black trees were invented by Bayer [3]. Guibas and Sedgwick [8] introduced the red/black color convention. Red-black trees can be seen as an encoding of 2-3-4 trees as binary trees.

Our starting point was an existing formalization in the Isabelle distribution (in `HOL/Library/RBT_Impl.thy`, by Reiter and Krauss) which in turn is based on the code by Okasaki [22] (for *insert*) and Stefan Kahrs [14] (for *delete* see the URL given in the article). The original verification has a certain similarity to ours because it also involves an inorder listing of the tree (function *entries*), but a number of the proofs are distinctly long and manual. In contrast, the only non-trivial lemmas we require are the following ones that need to be proved simultaneously about three auxiliary functions:

$$\begin{aligned} \uparrow [t] &\implies [del\ x\ t] = del_{list}\ x\ [t] \\ \uparrow [l] &\implies [delL\ x\ l\ a\ r] = del_{list}\ x\ [l] @ a \cdot [r] \\ \uparrow [r] &\implies [delR\ x\ l\ a\ r] = [l] @ a \cdot del_{list}\ x\ [r] \end{aligned}$$

Of course the proof is automatic, as usual.

Functional correctness of red-black trees can be proved without assuming any structural (red-black) invariants.

Related Work. Filliâtre and Letouzey [6] and Appel [2] verified red-black trees in Coq.

5.4 2-3 Trees

In a 2-3 tree (invented by Hopcroft in 1970 [5]), every non-leaf node has either two or three children: $\langle l, a, r \rangle$ or $\langle l, a, m, b, r \rangle$ where l, m, r are trees and a, b are elements. One can view $\langle l, a, m, b, r \rangle$ as a more compact representation of $\langle l, a, \langle m, b, r \rangle \rangle$ (see AA trees). Their structural invariant is that they are balanced, i.e. all leaves occur at the same depth.

Our code is based on the lecture notes by Turbak [29], who presents the key transformations in a graphical format. We present the more complex *delete* function in Fig. 4. Function *del* descends into the tree until the element (or a leaf) is found. Modified subtrees are recombined with smart constructors *nodeij* that combines i subtrees where subtree j has been modified and is wrapped up in either T_d (if the height of the subtree is unchanged) or Up_d (if the height of the subtree has decreased). We only show the functions *nodei1* because the other *nodeij* are symmetric.

The lemmas required for the correctness proof are similar to what we have seen already, with one new complication: the balancedness invariant *bal* is frequently required as a precondition, e.g. here:

$$del_min\ t = (x, t') \wedge bal\ t \wedge 0 < height\ t \implies x \cdot [tree_d\ t'] = [t]$$

Our automatic framework can cope because *bal* and *height* are defined in a straightforward manner by primitive recursion.

Related Work. The existing formalization of 2-3 trees in the Isabelle distribution (in `HOL/ex/Tree23.thy`, by Huffman and Nipkow) proves invariants but not functional correctness. Hoffmann and O'Donnell [12] give an equational definition of insertion. Reade [27] gives a similar equational definition of insertion and adds deletion; Turbak's version of deletion appears a bit simpler. Reade sketches (because there are too many cases) a pen-and-paper correctness proof and writes: "Mechanical support for such reasoning and the potential for partial automation of similar proofs are topics currently being investigated by the author".

5.5 2-3-4 Trees

2-3-4 trees are an extension of 2-3 trees where nodes may also have 4 children: $\langle t_1, a, t_2, b, t_3, c, t_4 \rangle$. Their structural invariant is that they are balanced, i.e. all leaves occur at the same depth. The code for 2-3-4 trees can also be viewed as an extension of that for 2-3 trees with additional cases. There are also new smart constructors *node4j*, e.g. *node41*:

$$\begin{aligned} node41\ (T_d\ t_1)\ a\ t_2\ b\ t_3\ c\ t_4 &= T_d\ \langle t_1, a, t_2, b, t_3, c, t_4 \rangle \\ node41\ (Up_d\ t_1)\ a\ \langle t_2, b, t_3 \rangle\ c\ t_4\ d\ t_5 &= T_d\ \langle \langle t_1, a, t_2, b, t_3 \rangle, c, t_4, d, t_5 \rangle \end{aligned}$$

```

datatype 'a upd = Td ('a tree23) | Upd ('a tree23)
treed (Td t) = t
treed (Upd t) = t

node21 (Td t1) a t2 = Td ⟨t1, a, t2⟩
node21 (Upd t1) a ⟨t2, b, t3⟩ = Upd ⟨t1, a, t2, b, t3⟩
node21 (Upd t1) a ⟨t2, b, t3, c, t4⟩ = Td ⟨⟨t1, a, t2⟩, b, ⟨t3, c, t4⟩⟩

node31 (Td t1) a t2 b t3 = Td ⟨t1, a, t2, b, t3⟩
node31 (Upd t1) a ⟨t2, b, t3⟩ c t4 = Td ⟨⟨t1, a, t2, b, t3⟩, c, t4⟩
node31 (Upd t1) a ⟨t2, b, t3, c, t4⟩ d t5 = Td ⟨⟨t1, a, t2⟩, b, ⟨t3, c, t4⟩, d, t5⟩

delmin ⟨⟨, a, ⟨⟩⟩ = (a, Upd ⟨⟩)
delmin ⟨⟨, a, ⟨⟩, b, ⟨⟩⟩ = (a, Td ⟨⟨, b, ⟨⟩⟩)
delmin ⟨l, a, r⟩ = (let (x, l') = delmin l in (x, node21 l' a r))
delmin ⟨l, a, m, b, r⟩ = (let (x, l') = delmin l in (x, node31 l' a m b r))

del x ⟨⟩ = Td ⟨⟩
del x ⟨⟨, a, ⟨⟩⟩ = (if x = a then Upd ⟨⟩ 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', t) = delmin r in node22 l a' t | 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') = delmin 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') = delmin r in node33 l a m b' r'
| GT ⇒ node33 l a m b (del x r))

delete x t = treed (del x t)

```

Fig. 4. Deletion in 2-3 trees

```

node41 (Upd t1) a ⟨t2, b, t3, c, t4⟩ d t5 e t6 =
Td ⟨⟨t1, a, t2⟩, b, ⟨t3, c, t4⟩, d, t5, e, t6⟩
node41 (Upd t1) a ⟨t2, b, t3, c, t4, d, t5⟩ e t6 f t7 =
Td ⟨⟨t1, a, t2⟩, b, ⟨t3, c, t4, d, t5⟩, e, t6, f, t7⟩

```

Related Work. It appears that the only (partially) published functional implementations of 2-3-4 trees is one in Maude [15] where the full code is available online. No formal proofs are reported.

5.6 1-2 Brother Trees

A 1-2 brother tree [23,24] is a binary tree with one further constructor M from trees to trees for unary nodes. The structural invariant is that the tree

is balanced (all leaves at the same depth) and that every unary node has a binary brother. Unary nodes allow us to balance any tree. There is a bijection between 1-2 brother trees and AVL trees: remove the unary nodes from a 1-2 brother tree and you obtain an AVL tree. Our formalization is based on the article by Hinze [10] where all code and invariants can be found. Hinze captures the invariant by two sets $B h$ and $U h$, the sets of brother trees of height h that have a binary (or nullary) respectively unary root node. The actual brother trees are captured by B ; U is an auxiliary notion. The correctness lemmas (1)–(3) for *insert*, *delete* and *isin* employ the abbreviation $T h = B h \cup U h$:

$$\begin{aligned} t \in T h \wedge \uparrow [t] &\Longrightarrow [insert\ a\ t] = ins_{list}\ a\ [t] \\ t \in T h \wedge \uparrow [t] &\Longrightarrow [delete\ x\ t] = del_{list}\ x\ [t] \\ t \in T h \wedge \uparrow [t] &\Longrightarrow isin\ t\ x = (x \in elems\ [t]) \end{aligned}$$

The non-trivial but automatic auxiliary lemmas are

$$\begin{aligned} t \in T h \wedge \uparrow [t] &\Longrightarrow [ins\ a\ t] = ins_{list}\ a\ [t] \\ t \in T h \wedge \uparrow [t] &\Longrightarrow [del\ x\ t] = del_{list}\ x\ [t] \\ t \in T h &\Longrightarrow \\ (del_min\ t = None) &= ([t] = []) \wedge \\ (del_min\ t = Some\ (a,\ t')) &\longrightarrow [t] = a \cdot [t'] \end{aligned}$$

5.7 AA Trees

Arne Anderson [1] invented a particularly simple form of balanced trees, named AA trees by Weiss [30]. They encode 2-3 trees as binary trees (with the help of an additional height field, although a single bit would suffice). Their main selling point is simplicity and compactness of the code. Our verification started from the functional version of AA trees published by Ragde [25] without proofs. The proofs for insertion were automatic as usual, but deletion posed problems.

The use of non-linear patterns in the Haskell code for `delete` was easily fixed. Then a failed correctness proof revealed that function `dellrg` goes down the wrong branch in the recursive case. After this bug was corrected the next complication was the fact that the definition of function `adjust` (which is supposed to restore the invariant after deletion) does not cover certain trees that cannot arise. Therefore I needed to introduce the following invariant corresponding to the textual invariants AA1–AA3 in [25]; function `lvl` returns the height field of a node:

$$\begin{aligned} invar\ \langle \rangle &= True \\ invar\ \langle h,\ l,\ a,\ r \rangle &= \\ (invar\ l \wedge invar\ r \wedge h &= lvl\ l + 1 \wedge \\ (h &= lvl\ r + 1 \vee (\exists\ lr\ b\ rr.\ r = \langle h,\ lr,\ b,\ rr \rangle \wedge h = lvl\ rr + 1))) \end{aligned}$$

Proving that insertion and deletion preserve the invariant was non-trivial, in particular because there were two more bugs:

- Function `dellrg` fails to call `adjust` to restore the invariant. This is the correct code (we call `dellrg` *del_max*):

$$\begin{aligned} del_max \langle lv, l, a, \langle \rangle \rangle &= (l, a) \\ del_max \langle lv, l, a, r \rangle &= (let (r', b) = del_max r in (adjust \langle lv, l, a, r' \rangle, b)) \end{aligned}$$

- The auxiliary function `nlvl` is incorrect. The correct version is as follows:
nlvl $t = (if\ snl\ t\ then\ lvl\ t\ else\ lvl\ t + 1)$

For the verification of functional correctness of deletion the domain of the partial `adjust` had to be characterized by a predicate *pre_adjust* (not in [25]). With its help we can formulate and prove the trivial *inorder*-lemma for *adjust*:

$$t \neq \langle \rangle \wedge pre_adjust\ t \implies \lfloor adjust\ t \rfloor = \lfloor t \rfloor$$

The main correctness theorem (2) requires a number of further lemmas:

$$\begin{aligned} del_max\ t &= (t', x) \wedge t \neq \langle \rangle \wedge invar\ t \implies \lfloor t' \rfloor @ [x] = \lfloor t \rfloor \\ invar\ \langle lv, l, a, r \rangle \wedge post_del\ l\ l' &\implies pre_adjust\ \langle lv, l', b, r \rangle \\ invar\ \langle lv, l, a, r \rangle \wedge post_del\ r\ r' &\implies pre_adjust\ \langle lv, l, a, r' \rangle \\ invar\ t \wedge (t', x) = del_max\ t \wedge t \neq \langle \rangle &\implies post_del\ t\ t' \\ invar\ t &\implies post_del\ t\ (delete\ x\ t) \end{aligned}$$

As usual, the proofs of the *inorder*-lemmas and theorems are automatic. The last four lemmas and the pre- and post-conditions involved are part of the invariant proofs and are merely reused. Hence they are not included in Table 1.

5.8 Splay Trees

Splay trees [28] are self-adjusting binary search trees where query and update operations modify the tree by rotating the accessed element to the root of the tree. The logarithmic amortized complexity of splay trees has been verified before [18]. The functional correctness proofs [17] followed the standard approach. Starting from the same code we automated those proofs.

Splay trees are different from the other trees we cover. All operations are based on a function *splay* :: 'a ⇒ 'a tree ⇒ 'a tree that rotates the given element (or an element close to it) to the root of the tree. For example, this is *isin*:

$$isin\ t\ x = (case\ splay\ x\ t\ of\ \langle \rangle \Rightarrow False \mid \langle l, a, r \rangle \Rightarrow x = a)$$

See elsewhere [17, 18] for *insert* and *delete*. Note that *isin* should return the new tree as well to achieve amortized logarithmic complexity. This is awkward in a functional language and gives the data structure an imperative flavour.

The verification is more demanding than before and we present all the required lemmas in Fig. 5. Lemmas (15)–(20) extend our lemma library in Fig. 3 but are only required for splay trees. With the help of these lemmas, the proofs of (1)–(3) are automatic.

$$\begin{aligned}
\uparrow (x \cdot xs) \wedge y \leq x &\Longrightarrow \uparrow (y \cdot xs) & (15) \\
\uparrow (xs @ [x]) \wedge x \leq y &\Longrightarrow \uparrow (xs @ [y]) & (16) \\
\uparrow (x \cdot xs) &\Longrightarrow \mathit{inslist} \ x \ xs = x \cdot xs & (17) \\
\uparrow (xs @ [x]) &\Longrightarrow \mathit{inslist} \ x \ xs = xs @ [x] & (18) \\
\uparrow (x \cdot xs) &\Longrightarrow \mathit{delist} \ x \ xs = xs & (19) \\
\uparrow (xs @ [x]) &\Longrightarrow \mathit{delist} \ x \ (xs @ ys) = xs @ \mathit{delist} \ x \ ys & (20) \\
(\mathit{splay} \ a \ t = \langle \rangle) &= (t = \langle \rangle) \\
(\mathit{splay_max} \ t = \langle \rangle) &= (t = \langle \rangle) \\
\mathit{splay} \ x \ t = \langle l, a, r \rangle \wedge \uparrow [t] &\Longrightarrow (x \in \mathit{elems} \ [t]) = (x = a) \\
[\mathit{splay} \ x \ t] &= [t] \\
\uparrow [t] \wedge \mathit{splay} \ x \ t = \langle l, a, r \rangle &\Longrightarrow \uparrow ([l] @ x \cdot [r]) \\
\mathit{splay_max} \ t = \langle l, a, r \rangle \wedge \uparrow [t] &\Longrightarrow [l] @ [a] = [t] \wedge r = \langle \rangle
\end{aligned}$$

Fig. 5. Lemmas for splay tree verification

6 Maps

6.1 Specifications

Search trees can implement maps as well as sets. Although sets are a special case of maps, we presented sets first because their simplicity facilitates the explanation of the basic concepts. Now we present the modifications required for maps. An implementation of maps must provide a type $(\prime a, \prime b) \ t$ (where $\prime a$ are the keys and $\prime b$ the values) with the operations

$$\begin{aligned}
\mathit{empty} &:: (\prime a, \prime b) \ t \\
\mathit{update} &:: \prime a \Rightarrow \prime b \Rightarrow (\prime a, \prime b) \ t \Rightarrow (\prime a, \prime b) \ t \\
\mathit{delete} &:: \prime a \Rightarrow (\prime a, \prime b) \ t \Rightarrow (\prime a, \prime b) \ t \\
\mathit{lookup} &:: (\prime a, \prime b) \ t \Rightarrow \prime a \Rightarrow \prime b \ \mathit{option}
\end{aligned}$$

where **datatype** $\prime a \ \mathit{option} = \mathit{None} \mid \mathit{Some} \ \prime a$ is predefined. Function lookup also plays the role of the abstraction function. In addition there is a data type invariant $\mathit{invar} :: (\prime a, \prime b) \ t \Rightarrow \mathit{bool}$. The specification of maps is shown in Fig. 6 (corresponding to Fig. 1). It uses the function update notation

$$f(a := b) = (\lambda x. \mathit{if} \ x = a \ \mathit{then} \ b \ \mathit{else} \ f \ x)$$

$$\begin{aligned}
\mathit{invar} \ m &\Longrightarrow \mathit{lookup} \ (\mathit{update} \ a \ b \ m) = (\mathit{lookup} \ m)(a := \mathit{Some} \ b) \\
\mathit{invar} \ m &\Longrightarrow \mathit{lookup} \ (\mathit{delete} \ a \ m) = (\mathit{lookup} \ m)(a := \mathit{None}) \\
\mathit{invar} \ m &\Longrightarrow \mathit{invar} \ (\mathit{update} \ a \ b \ m) \\
\mathit{invar} \ m &\Longrightarrow \mathit{invar} \ (\mathit{delete} \ a \ m)
\end{aligned}$$

Fig. 6. Specification of map implementations

Now we assume that the keys are linearly ordered. Search trees are abstracted to a list of key-value pairs sorted by their keys. The auxiliary functions \uparrow , $\mathit{inslist}$ and delist are replaced by

$$\begin{aligned}
& \uparrow_1 :: ('a \times 'b) \text{ list} \Rightarrow \text{bool} \\
& \text{upd}_{list} :: 'a \Rightarrow 'b \Rightarrow ('a \times 'b) \text{ list} \Rightarrow ('a \times 'b) \text{ list} \\
& \text{del}_{list} :: 'a \Rightarrow ('a \times 'b) \text{ list} \Rightarrow ('a \times 'b) \text{ list} \\
- & \uparrow_1 \text{ xs} = \uparrow (\text{map fst xs}) \text{ where } \text{fst } (a, b) = a. \\
- & \text{upd}_{list} \ a \ b \text{ updates a sorted (w.r.t. } \uparrow_1 \text{) list by either inserting } (a, b) \text{ at the} \\
& \text{correct position (w.r.t. } < \text{) if no } (a, -) \text{ is in the list, or replacing the first } (a, -) \\
& \text{by } (a, b) \text{ otherwise.} \\
- & \text{del}_{list} \ a \text{ deletes the first occurrence of a pair } (a, -) \text{ from a list.} \\
& \text{invar } t \Longrightarrow [\text{update } a \ b \ t] = \text{upd}_{list} \ a \ b \ [t] \\
& \text{invar } t \Longrightarrow [\text{delete } a \ t] = \text{del}_{list} \ a \ [t] \\
& \text{invar } t \Longrightarrow \text{lookup } t \ a = \text{map_of } [t] \ a \\
& \text{invar } t \Longrightarrow \text{inv } (\text{update } a \ b \ t) \\
& \text{invar } t \Longrightarrow \text{inv } (\text{delete } a \ t)
\end{aligned}$$

Fig. 7. Specification of map implementations over ordered types

Our second specification of maps (over a linearly ordered type $'a$) is shown in Fig. 7 (corresponding to Fig. 2). It is again based on an inorder function:

$$\text{inorder} :: ('a, 'b) \ t \Rightarrow ('a \times 'b) \ \text{list}$$

Again, we abbreviate $\text{inorder } t$ by $[t]$.

The search tree invariant is now expressed as $\uparrow_1 [t]$. Structural invariants can be added via the specification function $\text{inv} :: ('a, 'b) \ t \Rightarrow \text{bool}$ and we define

$$\text{invar } t = (\text{inv } t \wedge \uparrow_1 [t])$$

The first three propositions in Fig. 7 express functional correctness of update , delete and lookup w.r.t. upd_{list} , del_{list} and map_of . The latter is a predefined function on key-value lists:

$$\begin{aligned}
\text{map_of } [] &= (\lambda_. \text{None}) \\
\text{map_of } ((a, b) \cdot ps) &= (\text{map_of } ps)(a := b)
\end{aligned}$$

The next two propositions demand that inv is invariant. It is easy to show that Fig. 7 implies Fig. 6.

6.2 Proof Automation

Figure 8 (corresponding to Fig. 3) shows the set of lemmas used to automate the correctness proofs of implementations of maps. There are no lemmas about \uparrow_1 because its definition is simply unfolded and the lemmas (6)–(9) about \uparrow apply.

The litmus tests for the lemma collection are the correctness proofs for the map-variants of all the search trees discussed in Sect. 5. The code of the map-variants is structurally the same as their set-counterparts. The same is true for the lemmas required in the verification. In the end, the proofs of the map-variants are just as automatic as the ones of their set-counterparts.

$$\begin{aligned}
& \uparrow_1 (ps \text{ @ } [(a, b)]) \implies upd_{list} \ x \ y \ (ps \text{ @ } (a, b) \cdot qs) = \\
& \quad (if \ x < a \ then \ upd_{list} \ x \ y \ ps \ \text{@} \ (a, b) \cdot qs \ else \ ps \ \text{@} \ upd_{list} \ x \ y \ ((a, b) \cdot qs)) \\
& \uparrow_1 (ps \ \text{@} \ (a, b) \cdot qs) \implies del_{list} \ x \ (ps \ \text{@} \ (a, b) \cdot qs) = \\
& \quad (if \ x < a \ then \ del_{list} \ x \ ps \ \text{@} \ (a, b) \cdot qs \ else \ ps \ \text{@} \ del_{list} \ x \ ((a, b) \cdot qs)) \\
& map_of \ (ps \ \text{@} \ qs) \ x = \\
& \quad (case \ map_of \ ps \ x \ of \ None \ \Rightarrow \ map_of \ qs \ x \ | \ Some \ y \ \Rightarrow \ Some \ y) \\
& \uparrow (a \cdot map \ fst \ ps) \wedge x < a \implies map_of \ ps \ x = None \\
& \uparrow (map \ fst \ ps \ \text{@} \ [a]) \wedge a \leq x \implies map_of \ ps \ x = None
\end{aligned}$$

Fig. 8. Lemmas for upd_{list} , del_{list} and map_of

7 Conclusion

Our proof method works well because all the trees we considered follow the same ordering principle: inorder traversal yields a sorted list. Two referees suspected that for Trie-like trees [7] it would not work so well. I formalized binary trees where nodes are addressed by bit lists indicating the path to the node. A direct correctness proof is easy. The methods of this paper can also be applied (the list of addresses of the nodes in a tree, in prefix order, is lexicographically ordered) but the proof is more complicated and less automatic. Our approach seems overkill and awkward for such search trees.

Acknowledgement. Daniel Stüwe found and corrected the two invariant-related bugs in AA trees and proved preservation of the invariant under deletion for AA trees and 1-2 Brother trees.

References

1. Andersson, A.: Balanced search trees made simple. In: Dehne, F., Sack, J.-R., Santoro, N. (eds.) WADS 1993. LNCS, vol. 709, pp. 60–71. Springer, Heidelberg (1993)
2. Appel, A.: Efficient verified red-black trees (2011)
3. Bayer, R.: Symmetric binary B-trees: Data structure and maintenance algorithms. *Acta Inform.* **1**, 290–306 (1972)
4. Clochard, M.: Automatically verified implementation of data structures based on AVL trees. In: Giannakopoulou, D., Kroening, D. (eds.) VSTTE 2014. LNCS, vol. 8471, pp. 167–180. Springer, Heidelberg (2014)
5. Cormen, T.H., Leiserson, C.E., Rivest, R.L.: *Introduction to Algorithms*. MIT Press, Cambridge (1990)
6. Filliâtre, J.-C., Letouzey, P.: Functors for proofs and programs. In: Schmidt, D. (ed.) ESOP 2004. LNCS, vol. 2986, pp. 370–384. Springer, Heidelberg (2004)
7. Fredkin, E.: Trie memory. *Commun. ACM* **3**(9), 490–499 (1960)
8. Guibas, L.J., Sedgwick, R.: A dichromatic framework for balanced trees. In: 19th Annual Symposium on Foundations of Computer Science, pp. 8–21. IEEE Computer Society (1978)

9. Haftmann, F., Krauss, A., Kunčar, O., Nipkow, T.: Data refinement in Isabelle/HOL. In: Blazy, S., Paulin-Mohring, C., Pichardie, D. (eds.) ITP 2013. LNCS, vol. 7998, pp. 100–115. Springer, Heidelberg (2013)
10. Hinze, R.: Purely functional 1-2 brother trees. *J. Funct. Program.* **19**(6), 633–644 (2009)
11. Hoare, C.: Proof of correctness of data representations. *Acta Inform.* **1**, 271–281 (1972)
12. Hoffmann, C.M., O’Donnell, M.J.: Programming with equations. *ACM Trans. Program. Lang. Syst.* **4**(1), 83–112 (1982)
13. Jones, C.B.: *Software Development. A Rigorous Approach*. Prentice Hall, London (1980)
14. Kahrs, S.: Red black trees with types. *J. Funct. Program.* **11**(4), 425–432 (2001)
15. Martí-Oliet, N., Palomino, M., Verdejo, A.: A tutorial on specifying data structures in Maude. *Electr. Notes Theor. Comput. Sci.* **137**(1), 105–132 (2005)
16. Nipkow, T.: Are homomorphisms sufficient for behavioural implementations of deterministic and nondeterministic data types? In: Brandenburg, F.J., Wirsing, M., Vidal-Naquet, G. (eds.) STACS 1987. LNCS, vol. 247, pp. 260–271. Springer, Heidelberg (1987)
17. Nipkow, T.: Splay tree. *Archive of Formal Proofs, Formal proof development*, August 2014. http://isa-afp.org/entries/Splay_Tree.shtml
18. Nipkow, T.: Amortized complexity verified. In: Urban, C., Zhang, X. (eds.) ITP 2015. LNCS, vol. 9236, pp. 310–324. Springer, Heidelberg (2015)
19. Nipkow, T., Klein, G.: *Concrete Semantics with Isabelle/HOL*. Springer (2014). <http://concrete-semantics.org>
20. Nipkow, T., Kunčar, O., Pusch, C.: AVL trees. *Archive of Formal Proofs, Formal proof development*, March 2004. <http://isa-afp.org/entries/AVL-Trees.shtml>
21. Nipkow, T., Paulson, L.C., Wenzel, M. (eds.): *Isabelle/HOL. A Proof Assistant for Higher-Order Logic*. LNCS, vol. 2283. Springer, Heidelberg (2002)
22. Okasaki, C.: *Purely Functional Data Structures*. Cambridge University Press, Cambridge (1998)
23. Ottmann, T., Six, H.W.: Eine neue Klasse von ausgeglichenen Binärbäumen. *Angeordnete Informatik* **18**(9), 395–400 (1976)
24. Ottmann, T., Wood, D.: 1-2 brother trees or AVL trees revisited. *Comput. J.* **23**(3), 248–255 (1980)
25. Ragde, P.: Simple balanced binary search trees. In: Caldwell, J., Hölzenspies, P., Achten, P. (eds.) *Trends in Functional Programming in Education*, EPTCS, vol. 170, pp. 78–87 (2014)
26. Ralston, R.: ACL2-certified AVL trees. In: *Proceedings of 8th International Workshop ACL2 Theorem Prover and its Applications*, pp. 71–74. ACM (2009)
27. Reade, C.: Balanced trees with removals: an exercise in rewriting and proof. *Sci. Comput. Program.* **18**(2), 181–204 (1992)
28. Sleator, D.D., Tarjan, R.E.: Self-adjusting binary search trees. *J. ACM* **32**(3), 652–686 (1985)
29. Turbak, F.: CS230 Handouts – Spring 2007 (2007). <http://cs.wellesley.edu/cs230/spring07/handouts.html>
30. Weiss, M.A.: *Data Structures and Algorithm Analysis*, 2nd edn. Benjamin/Cummings, Redwood City (1994)