Amortized Complexity Verified

Tobias Nipkow
Technische Universität München

Abstract

A framework for the analysis of the amortized complexity of (functional) data structures is formalized in Isabelle/HOL and applied to a number of standard examples and to three famous non-trivial ones: skew heaps, splay trees and splay heaps.

1 Introduction

Amortized complexity [14, 3] of an algorithm averages the running times of a sequence of invocations of the algorithm. In this paper we formalize a simple framework for the analysis of the amortized complexity of functional programs and apply it to both the easy standard examples and the more challenging examples of skew heaps, splay trees and splay heaps. We have also analyzed pairing heaps [4] but cannot present them here for lack of space. All proofs are available online [9].

We are aiming for a particularly lightweight framework that supports proofs at a high level of abstraction. Therefore all algorithms are modeled as recursive functions in the logic. Because mathematical functions do not have a complexity they need to be accompanied by definitions of the intended running time that follow the recursive structure of the actual function definitions. Thus the user is free to choose the level of granularity of the complexity analysis. In our examples we simply count function calls.

Although one can view the artefacts that we analyze as functional programs, one can also view them as models or abstractions of imperative programs. For the amortized complexity we are only interested in the input-output behaviour and the running time complexity. As long as those are the same, it does not matter in what language the algorithm is implemented. In fact, the standard imperative implementations of all of our examples have the same complexity as the functional model. However, in a functional setting, amortized complexity reasoning may be invalid if the data structure under consideration is not used in a single-threaded manner [10].

1.1 Related Work

Hofmann and Jost [7] pioneered automatic type-based amortized analysis of heap usage of functional programs. This was later generalized in many directions, for example allowing multivariate polynomials [6]. Atkey [1] carries some of the ideas over to an imperative language with separation logic embedded in Coq. Charguéraud and Pottier [2] employ separation logic to verify the almost-linear amortized complexity of a Union-Find implementation in OCaml in Coq.

2 Lists and Trees

Lists are constructed from the empty list $[]$ via the infix cons-operator "$\cdot$", $|xs|$ is the length of xs, tl takes the tail and rev reverses a list.

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 $a :: 'a$. The size of a tree is the number of its nodes:

$$|\langle\rangle| = 0 \qquad |\langle l,\ \_,\ r\rangle| = |l| + |r| + 1$$

For convenience there is also the modified size function $|t|_1 = |t| + 1$.

3 Amortized Analysis Formalized

We formalize the following scenario. We are given a number of operations that may take parameters and that update the state (some data structure) as a side effect. Our model is purely functional: the state is replaced by a new state with each invocation of an operation, rather than mutating the state. This makes no difference because we only analyze time, not space.

Our model of amortized analysis is a theory that is parameterized as follows (a *locale* in Isabelle-speak):

$'s$ is the type of state.

$'o$ is the type of operations.

$init :: 's$ is the initial state.

$nxt :: 'o \Rightarrow 's \Rightarrow 's$ is the next state function.

$inv :: 's \Rightarrow bool$ is an invariant on the state. We assume that the invariant holds initially ($inv\ init$) and that it is preserved by all operations ($inv\ s \implies inv\ (nxt\ f\ s)$).

$t :: 'o \Rightarrow 's \Rightarrow real$ is the timing function: $t\ f\ s$ represents the time it takes to execute operation f in state s, i.e. $nxt\ f\ s$.

The effect of each operation f is modeled as a function $nxt\ f$ from state to state. Since functions are extensional, the execution time is modeled explicitly by function t. Alternatively one can instrument nxt with timing information and have it return a pair of a new state and the time the operation took. We have separated the computation of the result and the timing information into two functions because that is what one would typically do in a first step anyway to simplify the proofs. In particular this means that t need not be a closed form expression for the actual complexity. In all of our examples the definition of t will follow the (usually recursive) structure of the definition of nxt precisely. One could go one step further and derive t from an intensional formulation of nxt automatically, but that is orthogonal to our aim in this paper, namely *amortized* complexity.

For the analysis of amortized complexity we formalize the *potential method*. That is, our theory has another parameter:

$\Phi :: \ 's \Rightarrow real$ is the *potential* of a state. We assume the potential is initially 0 ($\Phi\ init = 0$) and never becomes negative ($inv\ s \implies 0 \leq \Phi\ s$).

The potential of the state represents the savings that can pay for future restructurings of the data structure. Typically, the higher the potential, the more out of balance the data structure is. Note that the potential is just a means to an end, the analysis, but does not influence the actual operations.

Let us now analyze the complexity of a sequence of operations formalized as a function of type $nat \Rightarrow 'o$. The sequence is infinite for convenience but of course we only analyze the execution of finite prefixes. For this purpose we define $state\ f\ n$, the state after the execution of the first n elements of f:

$state :: (nat \Rightarrow 'o) \Rightarrow nat \Rightarrow 's$
$state\ f\ 0 = init$
$state\ f\ (Suc\ n) = nxt\ (f\ n)\ (state\ f\ n)$

Now we can define the amortized complexity of an operation as the actual time it takes plus the difference in potential:

$a :: (nat \Rightarrow 'o) \Rightarrow nat \Rightarrow real$
$a\ f\ i = t\ (f\ i)\ (state\ f\ i) + \Phi\ (state\ f\ (i+1)) - \Phi\ (state\ f\ i)$

By telescoping (i.e. induction) we obtain

$(\sum_{i<n}\ t\ (f\ i)\ (state\ f\ i)) = (\sum_{i<n}\ a\ f\ i) + \Phi\ (state\ f\ 0) - \Phi\ (state\ f\ n)$

where $\sum_{i<n} F\ i$ is the sum of all $F\ i$ with $i < n$. Because of the assumptions on Φ this implies that on average the amortized complexity is an upper bound of the real complexity:

$(\sum_{i<n}\ t\ (f\ i)\ (state\ f\ i)) \leq (\sum_{i<n}\ a\ f\ i)$

To complete our formalization we add one more parameter:

$U :: {}'o \Rightarrow {}'s \Rightarrow real$ is an explicit upper bound for the amortized complexity of each operation (in a certain state), i.e. we assume that

$$inv\ s \implies t\ f\ s\ +\ \Phi\ (nxt\ f\ s)\ -\ \Phi\ s\ \leq\ U\ f\ s$$

Thus we obtain that U is indeed an upper bound of the real complexity:

$$(\textstyle\sum_{i<n}\ t\ (f\ i)\ (state\ f\ i)) \leq (\textstyle\sum_{i<n}\ U\ (f\ i)\ (state\ f\ i))$$

Instantiating this theory of amortized complexity means defining the parameters and proving the assumptions, in particular about U.

4   Easy Examples

Unless noted otherwise, the examples in this section come from a standard textbook [3].

4.1   Binary Counter

We start with the binary counter explained in the introduction. The state space ${}'s$ is just a list of booleans, starting with the least significant bit. There is just one parameterless operation "increment". Thus we can model type ${}'o$ with the unit type. The increment operation is defined recursively:

$$incr\ [] = [True]$$
$$incr\ (False \cdot bs) = True \cdot bs$$
$$incr\ (True \cdot bs) = False \cdot incr\ bs$$

In complete analogy the running time function for $incr$ is defined:

$$t_{incr}\ [] = 1$$
$$t_{incr}\ (False \cdot bs) = 1$$
$$t_{incr}\ (True \cdot bs) = t_{incr}\ bs\ +\ 1$$

Now we can instantiate the parameters of the amortized analysis theory:

$$init = []\qquad nxt\ () = incr\qquad t\ () = t_{incr}$$
$$inv\ s = True\qquad \Phi\ s = |filter\ id\ s|\qquad U\ ()\ s = 2$$

The key idea of the analysis is to define the potential of s as $|filter\ id\ s|$, the number of $True$ bits in s. This makes sense because the higher the potential, the longer an increment may take (roughly speaking). Now it is easy to show that 2 is an upper bound for the amortized complexity: the requirement on U follows immediately from this lemma (which is proved by induction):

$$t_{incr}\ bs\ +\ \Phi_{incr}\ (incr\ bs)\ -\ \Phi_{incr}\ bs = 2$$

4

4.2 Stack with Multipop

The operations are

datatype $'a\ op_{stk} = Push\ 'a \mid Pop\ nat$

where $Pop\ n$ pops n elements off the stack:

$nxt_{stk}\ (Push\ x)\ xs = x \cdot xs$
$nxt_{stk}\ (Pop\ n)\ xs = drop\ n\ xs$

In complete analogy the running time function is defined:

$t_{stk}\ (Push\ x)\ xs = 1$
$t_{stk}\ (Pop\ n)\ xs = min\ n\ |xs|$

Now we can instantiate the parameters of the amortized analysis theory:

$init = []$ $\quad\quad nxt = nxt_{stk}$ $\quad t = t_{stk}$ $\quad inv\ s = True$
$\Phi = length$ $\quad\quad U\ f\ s = (\text{case } f \text{ of } Push\ x \Rightarrow 2 \mid Pop\ n \Rightarrow 0)$

The necessary proofs are all automatic.

4.3 Dynamic Tables

Dynamic tables are tables where elements are added and deleted and the table grows and shrinks accordingly. We ignore the actual elements because they are irrelevant for the complexity analysis. Therefore the operations

datatype $op_{tb} = Ins \mid Del$

do not have arguments. Similarly the state is merely a pair of natural numbers (n, l) that abstracts a table of size l with n elements. This is how the operations behave:

$nxt_{tb}\ Ins\ (n,\ l) = (n + 1,\ \text{if } n < l \text{ then } l \text{ else if } l = 0 \text{ then } 1 \text{ else } 2 * l)$
$nxt_{tb}\ Del\ (n,\ l) =$
$(n - 1,\ \text{if } n = 1 \text{ then } 0 \text{ else if } 4 * (n - 1) < l \text{ then } l \text{ div } 2 \text{ else } l)$

If the table overflows upon insertion, its size is doubled. If a table is less than one quarter full after deletion, its size is halved. The transition from and to the empty table is treated specially.

This is the corresponding running time function:

$t_{tb}\ Ins\ (n,\ l) = (\text{if } n < l \text{ then } 1 \text{ else } n + 1)$
$t_{tb}\ Del\ (n,\ l) = (\text{if } n = 1 \text{ then } 1 \text{ else if } 4 * (n - 1) < l \text{ then } n \text{ else } 1)$

The running time for the cases where the table expands or shrinks is determined by the number of elements that need to be copied.

Now we can instantiate the parameters of the amortized analysis theory. We start with the system itself:

$$init = (0, 0) \qquad nxt = nxt_{tb} \qquad t = t_{tb}$$
$$inv\ (n,l) = (\text{if } l = 0 \text{ then } n = 0 \text{ else } n \le l \wedge l \le 4 * n)$$

This is the first time we have a non-trivial invariant. The potential is also more complicated than before:

$$\Phi\ (n,l) = (\text{if } 2 * n < l \text{ then } l\ /\ 2 - n \text{ else } 2 * n - l)$$

Now it is automatic to show the following amortized complexity:

$$U\ f\ s = (\text{case } f \text{ of } Ins \Rightarrow 3 \mid Del \Rightarrow 2)$$

4.4 Queues

Queues have one operation for enqueueing a new item and one for dequeueing the oldest item:

datatype $'a\ op_q = Enq\ 'a \mid Deq$

We ignore accessing the oldest item because it is a constant time operation in our implementation.

The simplest possible implementation of functional queues (e.g. [10]) consists of two lists (stacks) (xs,ys):

$$nxt_q\ (Enq\ x)\ (xs,\ ys) = (x \cdot xs,\ ys)$$
$$nxt_q\ Deq\ (xs,\ ys) = (\text{if } ys = [] \text{ then } ([],\ tl\ (rev\ xs)) \text{ else } (xs,\ tl\ ys))$$

$$t_q\ (Enq\ x)\ (xs,\ ys) = 1$$
$$t_q\ Deq\ (xs,\ ys) = (\text{if } ys = [] \text{ then } |xs| \text{ else } 0)$$

Note that the running time function counts only allocations of list cells and that it assumes rev is linear. Now we can instantiate the parameters of the amortized analysis theory to show that the average complexity of both Enq and Deq is 2. The necessary proofs are all automatic.

$$init = ([],\ []) \qquad nxt = nxt_q \qquad t = t_q \qquad inv\ s = True$$
$$\Phi\ (xs,ys) = |xs| \qquad U\ f\ s = (\text{case } f \text{ of } Enq\ x \Rightarrow 2 \mid Deq \Rightarrow 0)$$

In the same manner I have also verified [9] a modified implementation where reversal happens already when $|xs| = |ys| + 1$; this improves the worst-case behaviour but (using $\Phi(xs,ys) = 2 * |xs|$) the amortized complexity of Enq increases to 3.

5 Skew Heaps

This section analyzes a beautifully simple data structure for priority queues: skew heaps [13]. Heaps are trees where the least element is at the root. We assume that the elements are linearly ordered. The central operation on skew heaps is *meld*, that merges two skew heaps and swaps children along the merge path:

$$meld \; h_1 \; h_2 =$$
$$(\textsf{case} \; h_1 \; \textsf{of} \; \langle\rangle \Rightarrow h_2$$
$$| \; \langle l_1, \; a_1, \; r_1\rangle \Rightarrow$$
$$\textsf{case} \; h_2 \; \textsf{of} \; \langle\rangle \Rightarrow h_1$$
$$| \; \langle l_2, \; a_2, \; r_2\rangle \Rightarrow$$
$$\textsf{if} \; a_1 \leq a_2 \; \textsf{then} \; \langle meld \; h_2 \; r_1, \; a_1, \; l_1\rangle \; \textsf{else} \; \langle meld \; h_1 \; r_2, \; a_2, \; l_2\rangle)$$

We consider the two operations of inserting an element and removing the minimal element:

datatype $'a \; op_{pq} = Insert \; 'a \; | \; Delmin$

They are implemented via *meld* as follows:

$$nxt_{pq} \; (Insert \; a) \; h = meld \; \langle\langle\rangle, \; a, \; \langle\rangle\rangle \; h$$
$$nxt_{pq} \; Delmin \; h = del\_min \; h$$

$$del\_min \; \langle\rangle = \langle\rangle$$
$$del\_min \; \langle l, \; m, \; r\rangle = meld \; l \; r$$

For the functional correctness proofs see [9].

The analysis by Sleator and Tarjan is not ideal as a starting point for a formalization. Luckily there is a nice, precise functional account by Kaldewaij and Schoenmakers [8] that we will follow (although their *meld* differs slightly from ours). Their cost measure counts the number of calls of *meld*, *Insert* and *Delmin*:

$$t_{meld} \; \langle\rangle \; h = 1$$
$$t_{meld} \; h \; \langle\rangle = 1$$
$$t_{meld} \; \langle l_1, \; a_1, \; r_1\rangle \; \langle l_2, \; a_2, \; r_2\rangle =$$
$$(\textsf{if} \; a_1 \leq a_2 \; \textsf{then} \; t_{meld} \; \langle l_2, \; a_2, \; r_2\rangle \; r_1 \; \textsf{else} \; t_{meld} \; \langle l_1, \; a_1, \; r_1\rangle \; r_2) + 1$$

$$t_{pq} \; (Insert \; a) \; h = t_{meld} \; \langle\langle\rangle, \; a, \; \langle\rangle\rangle \; h + 1$$
$$t_{pq} \; Delmin \; h = (\textsf{case} \; h \; \textsf{of} \; \langle\rangle \Rightarrow 1 \; | \; \langle t_1, \; a, \; t_2\rangle \Rightarrow t_{meld} \; t_1 \; t_2 + 1)$$

Kaldewaij and Schoenmakers prove a tighter upper bound than Sleator and Tarjan, replacing the factor of 3 by 1.44. We are satisfied with verifying the bound by Sleator and Tarjan and work with the following simple potential function which is an instance of the one by Kaldewaij and Schoenmakers: it counts the number of "right heavy" nodes.

$$\Phi \; \langle\rangle = 0$$
$$\Phi \; \langle l, \; \_, \; r\rangle = \Phi \; l + \Phi \; r + (\textsf{if} \; |l| < |r| \; \textsf{then} \; 1 \; \textsf{else} \; 0)$$

To prove the amortized complexity of *meld* we need some further notions that capture the ideas of Sleator and Tarjan in a concise manner:

$$rheavy \; \langle l, \, \_, \, r \rangle = (|l| < |r|)$$

$$lpath \; \langle \rangle = []$$
$$lpath \; \langle l, \, a, \, r \rangle = \langle l, \, a, \, r \rangle \cdot lpath \; l$$

$$rpath \; \langle \rangle = []$$
$$rpath \; \langle l, \, a, \, r \rangle = \langle l, \, a, \, r \rangle \cdot rpath \; r$$

$$\Gamma \; h = |filter \; rheavy \; (lpath \; h)|$$
$$\Delta \; h = |filter \; (\lambda p. \; \neg \; rheavy \; p) \; (rpath \; h)|$$

Two easy inductive properties:

$$\Gamma \; h \le \log_2 |h|_1 \quad (1) \qquad \Delta \; h \le \log_2 |h|_1 \quad (2)$$

Now the desired logarithmic amortized complexity of *meld* follows:

$$t_{meld} \; t_1 \; t_2 + \Phi \; (meld \; t_1 \; t_2) - \Phi \; t_1 - \Phi \; t_2$$
$$\le \Gamma \; (meld \; t_1 \; t_2) + \Delta \; t_1 + \Delta \; t_2 + 1 \quad \text{by induction on } meld$$
$$\le \log_2 |meld \; t_1 \; t_2|_1 + \log_2 |t_1|_1 + \log_2 |t_2|_1 + 1 \quad \text{by (1), (2)}$$
$$= \log_2 (|t_1|_1 + |t_2|_1 - 1) + \log_2 |t_1|_1 + \log_2 |t_2|_1 + 1$$
$$\quad \text{because } |meld \; t_1 \; t_2| = |t_1| + |t_2|$$
$$\le \log_2 (|t_1|_1 + |t_2|_1) + \log_2 |t_1|_1 + \log_2 |t_2|_1 + 1$$
$$\le \log_2 (|t_1|_1 + |t_2|_1) + 2 * \log_2 (|t_1|_1 + |t_2|_1) + 1$$
$$\quad \text{because } \log_2 x + \log_2 y \le 2 * \log_2 (x + y) \text{ if } x, y \ge 1$$
$$= 3 * \log_2 (|t_1|_1 + |t_2|_1) + 1$$

Now it is easy to verify the following amortized complexity for *Insert* and *Delmin* by instantiating our standard theory with

$$U \; (Insert \; \_) \; h = 3 * \log_2 (|h|_1 + 2) + 2$$
$$U \; Delmin = 3 * \log_2 (|h|_1 + 2) + 4$$

Note that Isabelle supports implicit coercions, in particular from *nat* to *real*, that are inserted automatically [15].

6 Splay Trees

A splay tree [12] is a subtle self-adjusting binary search tree. It achieves its amortized logarithmic complexity by local rotations of subtrees along the access path. Its central operation is *splay* of type $'a \Rightarrow 'a \; tree \Rightarrow 'a \; tree$ that rotates the given element (of a linearly ordered type $'a$) to the root of the tree. Most presentations of *splay* confine themselves to this case where the given element is in the tree. If the given element is not in the tree, the last element found before a $\langle \rangle$ was met is rotated to the root. The complete definition is shown in Figure 1.

Given splaying, searching for an element in the tree is trivial: you splay with the given element and check if it ends up at the root. For insertion

$splay\ a\ \langle\rangle = \langle\rangle$
$splay\ a\ \langle cl,\ c,\ cr\rangle =$
($if\ a = c\ then\ \langle cl,\ c,\ cr\rangle$
 $else\ if\ a < c$
      $then\ case\ cl\ of\ \langle\rangle \Rightarrow \langle cl,\ c,\ cr\rangle$
        $|\ \langle bl,\ b,\ br\rangle \Rightarrow$
          $if\ a = b\ then\ \langle bl,\ a,\ \langle br,\ c,\ cr\rangle\rangle$
          $else\ if\ a < b$
            $then\ if\ bl = \langle\rangle\ then\ \langle bl,\ b,\ \langle br,\ c,\ cr\rangle\rangle$
              $else\ case\ splay\ a\ bl\ of$
                $\langle al,\ a',\ ar\rangle \Rightarrow \langle al,\ a',\ \langle ar,\ b,\ \langle br,\ c,\ cr\rangle\rangle\rangle$
            $else\ if\ br = \langle\rangle\ then\ \langle bl,\ b,\ \langle br,\ c,\ cr\rangle\rangle$
              $else\ case\ splay\ a\ br\ of$
                $\langle al,\ a',\ ar\rangle \Rightarrow \langle\langle bl,\ b,\ al\rangle,\ a',\ \langle ar,\ c,\ cr\rangle\rangle$
     $else\ case\ cr\ of\ \langle\rangle \Rightarrow \langle cl,\ c,\ cr\rangle$
        $|\ \langle bl,\ b,\ br\rangle \Rightarrow$
          $if\ a = b\ then\ \langle\langle cl,\ c,\ bl\rangle,\ a,\ br\rangle$
          $else\ if\ a < b$
            $then\ if\ bl = \langle\rangle\ then\ \langle\langle cl,\ c,\ bl\rangle,\ b,\ br\rangle$
              $else\ case\ splay\ a\ bl\ of$
                $\langle al,\ a',\ ar\rangle \Rightarrow \langle\langle cl,\ c,\ al\rangle,\ a',\ \langle ar,\ b,\ br\rangle\rangle$
            $else\ if\ br = \langle\rangle\ then\ \langle\langle cl,\ c,\ bl\rangle,\ b,\ br\rangle$
              $else\ case\ splay\ a\ br\ of$
                $\langle al,\ a',\ ar\rangle \Rightarrow \langle\langle\langle cl,\ c,\ bl\rangle,\ b,\ al\rangle,\ a',\ ar\rangle)$

Figure 1: Function $splay$

$insert_{st}\ a\ t =$
($if\ t = \langle\rangle\ then\ \langle\langle\rangle,\ a,\ \langle\rangle\rangle$
 $else\ case\ splay\ a\ t\ of$
     $\langle l,\ a',\ r\rangle \Rightarrow$
      $if\ a = a'\ then\ \langle l,\ a,\ r\rangle$
      $else\ if\ a < a'\ then\ \langle l,\ a,\ \langle\langle\rangle,\ a',\ r\rangle\rangle\ else\ \langle\langle l,\ a',\ \langle\rangle\rangle,\ a,\ r\rangle)$

$delete_{st}\ a\ t =$
($if\ t = \langle\rangle\ then\ \langle\rangle$
 $else\ case\ splay\ a\ t\ of$
     $\langle l,\ a',\ r\rangle \Rightarrow$
      $if\ a = a'$
      $then\ if\ l = \langle\rangle\ then\ r\ else\ case\ splay\_max\ l\ of\ \langle l',\ m,\ r'\rangle \Rightarrow \langle l',\ m,\ r\rangle$
      $else\ \langle l,\ a',\ r\rangle)$

Figure 2: Functions $insert_{st}$ and $delete_{st}$

$splay\_max\ \langle\rangle = \langle\rangle$
$splay\_max\ \langle l,\ b,\ \langle\rangle\rangle = \langle l,\ b,\ \langle\rangle\rangle$
$splay\_max\ \langle l,\ b,\ \langle rl,\ c,\ rr\rangle\rangle =$
$(\textbf{if}\ rr = \langle\rangle\ \textbf{then}\ \langle\langle l,\ b,\ rl\rangle,\ c,\ \langle\rangle\rangle$
$\ \textbf{else case}\ splay\_max\ rr\ \textbf{of}\ \langle rrl,\ x,\ xa\rangle \Rightarrow \langle\langle\langle l,\ b,\ rl\rangle,\ c,\ rrl\rangle,\ x,\ xa\rangle)$

Figure 3: Function $splay\_max$

$set_{tree}\ \langle\rangle = \emptyset$
$set_{tree}\ \langle l,\ a,\ r\rangle = \{a\} \cup (set_{tree}\ l \cup set_{tree}\ r)$

$bst\ \langle\rangle = True$
$bst\ \langle l,\ a,\ r\rangle =$
$(bst\ l \wedge bst\ r \wedge (\forall x \in set_{tree}\ l.\ x < a)\ \wedge (\forall x \in set_{tree}\ r.\ a < x))$

Figure 4: Functions set_{tree} and bst

and deletion, algorithm texts typically show pictures only. In contrast, we show the code only, in figures 2–3. To insert a, you splay with a to see if it is already there, and if it is not, you insert it at the top (which is the right place due to the previous splay action). To delete a, you splay with a and if a ends up at the root, you replace it with the maximal element removed from the left subtree. The latter step is performed by $splay\_max$ that splays with the maximal element.

6.1 Functional Correctness

So far we had ignored functional correctness but for splay trees we actually need it in the verification of the complexity. To formulate functional correctness we need the two auxiliary functions shown in Figure 4. Function set_{tree} collects the elements in the tree, function bst checks if the tree is a *binary search tree* according to the linear ordering "<" on the elements. The key functional properties are that splaying does not change the contents of the tree (it merely reorganizes it) and that bst is an invariant of splaying:

$set_{tree}\ (splay\ a\ t) = set_{tree}\ t$
$bst\ t \Longrightarrow bst\ (splay\ a\ t)$

Similar properties can be proved for insertion and deletion, e.g.,

$bst\ t \Longrightarrow set_{tree}\ (delete_{st}\ a\ t) = set_{tree}\ t - \{a\}$

Now we present two amortized analyses: a simpler one that yields the bounds proved by Sleator and Tarjan [12] and a more complicated and precise one due to Schoenmakers [11].
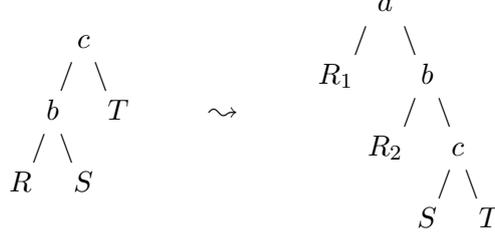
Figure 5: Zig-zig case for *splay*: $a < b < c$

6.2 Amortized Analysis

The timing functions are straightforward and not shown. Roughly speaking, they count only the number of splay steps: t_{splay} counts the number of calls of *splay*, t_{splay_max} counts the number of calls of *splay_max*; t_{delete} counts the time for both *splay* and *splay_max*.

The potential of a tree is defined as a sum of logarithms as follows:

$$\varphi\ t = \log_2 |t|_1$$
$$\Phi\ \langle\rangle = 0$$
$$\Phi\ \langle l,\ a,\ r\rangle = \Phi\ l + \Phi\ r + \varphi\ \langle l,\ a,\ r\rangle$$

The amortized complexity of splaying is defined as usual:

$$\mathcal{A}\ a\ t = t_{splay}\ a\ t + \Phi\ (splay\ a\ t) - \Phi\ t$$

Let *subtrees* yield the set of all subtrees of a tree:

$$subtrees\ \langle\rangle = \{\langle\rangle\}$$
$$subtrees\ \langle l,\ a,\ r\rangle = \{\langle l,\ a,\ r\rangle\} \cup (subtrees\ l \cup subtrees\ r)$$

The following logarithmic bound is proved by induction on t according to the recursion schema of *splay*: if *bst t* and $\langle l,\ a,\ r\rangle \in subtrees\ t$ then

$$\mathcal{A}\ a\ t \le 3 * (\varphi\ t - \varphi\ \langle l,\ a,\ r\rangle) + 1 \tag{3}$$

Let us look at one case of the inductive proof in detail. We pick the so-called zig-zig case shown in Figure 5. Subtrees with root x are called X on the left and X' on the right-hand side. Thus the figure depicts *splay a C = A'* assuming the recursive call *splay a R* = $\langle R_1,\ a,\ R_2\rangle =: R'$.

$$\mathcal{A}\ a\ C = \mathcal{A}\ a\ R + \varphi\ B' + \varphi\ C' - \varphi\ B - \varphi\ R' + 1$$
$$\le 3 * (\varphi\ R - \varphi\ \langle l,\ a,\ r\rangle) + \varphi\ B' + \varphi\ C' - \varphi\ B - \varphi\ R' + 2$$
$$\quad \text{by ind.hyp.}$$
$$= 2 * \varphi\ R + \varphi\ B' + \varphi\ C' - \varphi\ B - 3 * \varphi\ \langle l,\ a,\ r\rangle + 2$$

$\quad$ because $\varphi\ R = \varphi\ R'$
$\leq \varphi\ R + \varphi\ B' + \varphi\ C' - 3 * \varphi\ \langle l,\ a,\ r\rangle + 2$
$\quad$ because $\varphi\ B < \varphi\ R$
$\leq \varphi\ B' + 2 * \varphi\ C - 3 * \varphi\ \langle l,\ a,\ r\rangle + 1$
$\quad$ because $1 + \log_2 x + \log_2 y < 2 * \log_2 (x + y)$ if $x, y > 0$
$\leq 3 * (\varphi\ C - \varphi\ \langle l,\ a,\ r\rangle) + 1 \qquad$ because $\varphi\ B' \leq \varphi\ C$

This looks similar to the proof by Sleator and Tarjan but is different: they consider one double rotation whereas we argue about the whole input-output relationship; also our *log* argument is simpler.

$\quad$ From (3) we obtain in the worst case $(l = r = \langle\rangle)$:

$\quad$ If $bst\ t$ and $a \in set_{tree}\ t$ then $\mathcal{A}\ a\ t \leq 3 * (\varphi\ t - 1) + 1$.

In the literature the case $a \notin set_{tree}\ t$ is treated informally by stating that it can be reduced to $a' \in set_{tree}\ t$: one could have called *splay* with some $a' \in set_{tree}\ t$ instead of a and the behaviour would have been the same. Formally we prove by induction that if $t \neq \langle\rangle$ and $bst\ t$ then

$\quad \exists\, a' \in set_{tree}\ t.\ splay\ a'\ t = splay\ a\ t \wedge t_{splay}\ a'\ t = t_{splay}\ a\ t$

This gives us an upper bound for all binary search trees:

$$bst\ t \implies \mathcal{A}\ a\ t \leq 3 * \varphi\ t + 1 \qquad\qquad (4)$$

The $\varphi\ t - 1$ was increased to $\varphi\ t$ because the former is negative if $t = \langle\rangle$.
$\quad$ We also need to determine the amortized complexity $\mathcal{A}m$ of *splay_max*

$$\mathcal{A}m\ t = t_{splay\_max}\ t + \Phi\ (splay\_max\ t) - \Phi\ t$$

A derivation similar to but simpler than the one for $\mathcal{A}$ yields the same upper bound: $bst\ t \implies \mathcal{A}m\ t \leq 3 * \varphi\ t + 1$.
$\quad$ Now we can apply our amortized analysis theory:

$\quad$ **datatype** $'a\ op_{st} = Splay\ 'a\ |\ Insert\ 'a\ |\ Delete\ 'a$

$\quad nxt_{st}\ (Splay\ a)\ t = splay\ a\ t \qquad t_{st}\ (Splay\ a)\ t = t_{splay}\ a\ t$
$\quad nxt_{st}\ (Insert\ a)\ t = insert_{st}\ a\ t \qquad t_{st}\ (Insert\ a)\ t = t_{splay}\ a\ t$
$\quad nxt_{st}\ (Delete\ a)\ t = delete_{st}\ a\ t \qquad t_{st}\ (Delete\ a)\ t = t_{delete}\ a\ t$

$\quad init = \langle\rangle \qquad nxt = nxt_{st} \qquad t = t_{st} \qquad inv = bst \qquad \Phi = \Phi$
$\quad U\ (Splay\ \_)\ t = 3 * \varphi\ t + 1$
$\quad U\ (Insert\ \_)\ t = 4 * \varphi\ t + 2$
$\quad U\ (Delete\ \_)\ t = 6 * \varphi\ t + 2$

The fact that the given U is indeed a correct upper bound follows from the upper bounds for $\mathcal{A}$ and $\mathcal{A}m$; for *Insert* and *Delete* the proof needs more case distinctions and log-manipulations.

6.3 Improved Amortized Analysis

This subsection follows the work of Schoenmakers [11] (except that he confines himself to *splay*) who improves upon the constants in the above analysis. His analysis is parameterized by two constants $\alpha > 1$ and β subject to three constraints where all the variables are assumed to be ≥ 1:

$$(x + y) * (y + z)^\beta \leq (x + y)^\beta * (x + y + z)$$

$$\alpha * (l' + r') * (lr + r)^\beta * (lr + r' + r)^\beta$$
$$\leq (l' + r')^\beta * (l' + lr + r')^\beta * (l' + lr + r' + r)$$

$$\alpha * (l' + r') * (l' + ll)^\beta * (r' + r)^\beta$$
$$\leq (l' + r')^\beta * (l' + ll + r')^\beta * (l' + ll + r' + r)$$

The following upper bound is again proved by induction but this time with the help of the above constraints: if *bst t* and $\langle l, a, r \rangle \in \textit{subtrees } t$ then

$$\mathcal{A} \ a \ t \leq \log_\alpha \left(|t|_1 \ / \ (|l|_1 + |r|_1) \right) + 1$$

From this we obtain the following main theorem just like before:

$$\mathcal{A} \ a \ t \leq \log_\alpha |t|_1 + 1$$

Now we instantiate the above abstract development with $\alpha = \sqrt[3]{4}$ and $\beta = 1/3$ (which includes proving the three constraints on α and β above) to obtain a bound for splaying that is only half as large as in (4):

$$\textit{bst } t \implies \mathcal{A}_{34} \ a \ t \leq 3 \ / \ 2 * \varphi \ t$$

The subscript $_{34}$ is our indication that we refer to the $\alpha = \sqrt[3]{4}$ and $\beta = 1/3$ instance. Schoenmakers additionally showed that this specific choice of α and β yields the minimal upper bound.

A similar but simpler development leads to the same bound for $\mathcal{A}m_{34}$ as for $\mathcal{A}_{34}$. Again we apply our amortized analysis theory to verify upper bounds for *Splay*, *Insert* and *Delete* that are also only half as large as before:

$$U \ (\textit{Splay} \ \_) \ t = 3 \ / \ 2 * \varphi \ t + 1$$
$$U \ (\textit{Insert} \ \_) \ t = 2 * \varphi \ t + 3 \ / \ 2$$
$$U \ (\textit{Delete} \ \_) \ t = 3 * \varphi \ t + 2$$

The proofs in this subsection require a lot of highly nonlinear arithmetic. Only some of the polynomial inequalities can be automated with Harrison's sum-of-squares method [5].

7 Splay Heaps

Splay heaps are another self-adjusting data structure and were invented by Okasaki [10]. Splay heaps are organized internally like splay trees but they implement a priority queue interface. When inserting an element x into a splay heap, the splay heap is first partitioned (by rotations, like *splay*) into two trees, one $\leq x$ and one $> x$, and x becomes the new root:

$$insert\ x\ h = (\text{let } (l,\ r) = partition\ x\ h\ \text{in } \langle l,\ x,\ r \rangle)$$

$partition\ p\ \langle \rangle = (\langle \rangle,\ \langle \rangle)$
$partition\ p\ \langle al,\ a,\ ar \rangle =$
$(\text{if } a \leq p$
$\text{ then case } ar\ \text{of } \langle \rangle \Rightarrow (\langle al,\ a,\ ar \rangle,\ \langle \rangle)$
$\qquad | \ \langle bl,\ b,\ br \rangle \Rightarrow$
$\qquad\qquad \text{if } b \leq p\ \text{then let } (pl,\ y) = partition\ p\ br\ \text{in } (\langle \langle al,\ a,\ bl \rangle,\ b,\ pl \rangle,\ y)$
$\qquad\qquad \text{else let } (pl,\ pr) = partition\ p\ bl\ \text{in } (\langle al,\ a,\ pl \rangle,\ \langle pr,\ b,\ br \rangle)$
$\text{ else case } al\ \text{of } \langle \rangle \Rightarrow (\langle \rangle,\ \langle al,\ a,\ ar \rangle)$
$\qquad | \ \langle bl,\ b,\ br \rangle \Rightarrow$
$\qquad\qquad \text{if } b \leq p\ \text{then let } (pl,\ pr) = partition\ p\ br\ \text{in } (\langle bl,\ b,\ pl \rangle,\ \langle pr,\ a,\ ar \rangle)$
$\qquad\qquad \text{else let } (pl,\ pr) = partition\ p\ bl\ \text{in } (pl,\ \langle pr,\ b,\ \langle br,\ a,\ ar \rangle \rangle))$

Function *del_min* removes the minimal element and is similar to *splay_max*:

$del\_min\ \langle \rangle = \langle \rangle$
$del\_min\ \langle \langle \rangle,\ uu,\ r \rangle = r$
$del\_min\ \langle \langle ll,\ a,\ lr \rangle,\ b,\ r \rangle =$
$(\text{if } ll = \langle \rangle\ \text{then } \langle lr,\ b,\ r \rangle\ \text{else } \langle del\_min\ ll,\ a,\ \langle lr,\ b,\ r \rangle \rangle)$

In contrast to search trees, priority queues may contain elements multiple times. Therefore splay heaps satisfy the weaker invariant bst_{eq}:

$bst_{eq}\ \langle \rangle = True$
$bst_{eq}\ \langle l,\ a,\ r \rangle =$
$(bst_{eq}\ l \wedge bst_{eq}\ r \wedge (\forall x \in set_{tree}\ l.\ x \leq a)\ \wedge (\forall x \in set_{tree}\ r.\ a \leq x))$

This is an invariant for both *partition* and *del_min*:

If $bst_{eq}\ t$ and $partition\ p\ t = (l,\ r)$ then $bst_{eq}\ \langle l,\ p,\ r \rangle$.
If $bst_{eq}\ t$ then $bst_{eq}\ (del\_min\ t)$.

For the functional correctness proofs see [9].

7.1 Amortized Analysis

Now we verify the amortized analysis due to Okasaki. The timing functions are straightforward and not shown: t_{part} and t_{dm} count the number of calls

$$
\begin{array}{ccc}
\begin{array}{c}
a \\
/\ \backslash \\
b\quad T \\
/\ \backslash \\
R\quad S
\end{array}
& \rightsquigarrow &
\left(
\begin{array}{cc}
b & a \\
/\ \backslash & /\ \backslash \\
R\quad S_1 & S_2\quad T
\end{array}
\right)
\end{array}
$$
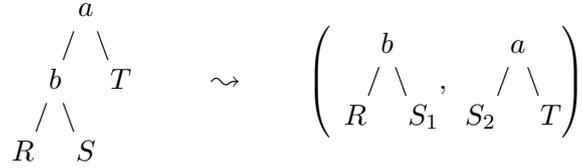
Figure 6: Zig-zag case for *partition*: $b \leq p < a$

of *partition* and *del_min*. The potential of a tree is defined as for splay trees in Section 6.2. The following logarithmic bound of the amortized complexity $\mathcal{A}\ p\ t = t_{part}\ p\ t + \Phi\ l' + \Phi\ r' - \Phi\ t$ is proved by computation induction on *partition* t: if $bst_{eq}\ t$ and *partition* $p\ t = (l', r')$ then

$$\mathcal{A}\ p\ t \leq 2 * \varphi\ t + 1$$

Okasaki [10] shows the zig-zig case of the induction, I show the zig-zag case in Figure 6. Subtrees with root x are called X on the left and X' on the right-hand side. Thus Figure 6 depicts *partition* $p\ A = (B', A')$ assuming the recursive call *partition* $p\ S = (S_1, S_2)$.

$$
\begin{aligned}
\mathcal{A}\ p\ A &= \mathcal{A}\ p\ S + 1 + \varphi\ B' + \varphi\ A' - \varphi\ B - \varphi\ A \\
&\leq 2 * \varphi\ S + 2 + \varphi\ B' + \varphi\ A' - \varphi\ B - \varphi\ A \qquad \text{by ind.hyp.} \\
&= 2 + \varphi\ B' + \varphi\ A' \qquad \text{because } \varphi\ S < \varphi\ B \text{ and } \varphi\ S < \varphi\ A \\
&\leq 2 * \log_2\ (|R|_1 + |S_1|_1 + |S_2|_1 + |T|_1 - 1) + 1 \\
&\quad \text{because } 1 + \log_2\ x + \log_2\ y \leq 2 * \log_2\ (x + y - 1) \text{ if } x, y \geq 2 \\
&= 2 * \varphi\ A + 1 \qquad \text{because } |S_1| + |S_2| = |S|
\end{aligned}
$$

The proof of the amortized complexity of *del_min* is similar to the one for *splay_max*: $t_{dm}\ t + \Phi\ (del\_min\ t) - \Phi\ t \leq 2 * \varphi\ t + 1$. Now it is routine to verify the following amortized complexities by instantiating our standard theory with $U\ (Insert\ \_)\ t = 3 * \log_2\ (|t|_1 + 1) + 1$ and $U\ Delmin\ t = 2 * \varphi\ t + 1$.

Acknowledgement Berry Schoenmakers patiently answered many questions about his work whenever I needed help.

References

[1] R. Atkey. Amortised resource analysis with separation logic. *Logical Methods in Computer Science*, 7(2), 2011.

[2] A. Charguéraud and F. Pottier. Machine-checked verification of the correctness and amortized complexity of an efficient union-find implementation. In C. Urban, editor, *ITP 2015*, volume ? of *LNCS*, pages ?–? Springer, 2015.

[3] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms.* MIT Press, 1990.

[4] M. L. Fredman, R. Sedgewick, D. D. Sleator, and R. E. Tarjan. The pairing heap: A new form of self-adjusting heap. *Algorithmica*, 1(1):111–129, 1986.

[5] J. Harrison. Verifying nonlinear real formulas via sums of squares. In K. Schneider and J. Brandt, editors, *TPHOLs 2007*, volume 4732 of *LNCS*, pages 102–118. Springer, 2007.

[6] J. Hoffmann, K. Aehlig, and M. Hofmann. Multivariate amortized resource analysis. *ACM Trans. Program. Lang. Syst.*, 34(3):14, 2012.

[7] M. Hofmann and S. Jost. Static prediction of heap space usage for first-order functional programs. In *Proc. 30th ACM Symposium Principles of Programming Languages*, pages 185–197, 2003.

[8] A. Kaldewaij and B. Schoenmakers. The derivation of a tighter bound for top-down skew heaps. *Information Processing Letters*, 37:265–271, 1991.

[9] T. Nipkow. Amortized complexity verified. *Archive of Formal Proofs*, 2014. http://afp.sf.net/entries/Amortized_Complexity.shtml, Formal proof development.

[10] C. Okasaki. *Purely Functional Data Structures.* Cambridge University Press, 1998.

[11] B. Schoenmakers. A systematic analysis of splaying. *Information Processing Letters*, 45:41–50, 1993.

[12] D. D. Sleator and R. E. Tarjan. Self-adjusting binary search trees. *J. ACM*, 32(3):652–686, 1985.

[13] D. D. Sleator and R. E. Tarjan. Self-adjusting heaps. *SIAM J. Comput.*, 15(1):52–69, 1986.

[14] R. E. Tarjan. Amortized complexity. *SIAM J. Alg. Disc. Meth.*, 6(2):306–318, 1985.

[15] D. Traytel, S. Berghofer, and T. Nipkow. Extending Hindley-Milner type inference with coercive structural subtyping. In H. Yang, editor, *APLAS 2011*, volume 7078 of *LNCS*, pages 89–104. Springer, 2011.