



The Functional Essence of Imperative Binary Search Trees

ANTON LORENZEN, University of Edinburgh, UK

DAAN LEIJEN, Microsoft Research, USA

WOUTER SWIERSTRA, Universiteit Utrecht, Netherlands

SAM LINDLEY, University of Edinburgh, UK

Algorithms on restructuring binary search trees are typically presented in imperative pseudocode. Understandably so, as their performance relies on in-place execution, rather than the repeated allocation of fresh nodes in memory. Unfortunately, these imperative algorithms are notoriously difficult to verify as their loop invariants must relate the unfinished tree fragments being rebalanced. This paper presents several novel functional algorithms for accessing and inserting elements in a restructuring binary search tree that are as fast as their imperative counterparts; yet the correctness of these functional algorithms is established using a simple inductive argument. For each data structure, move-to-root, splay, and zip trees, this paper describes both a bottom-up algorithm using zippers and a top-down algorithm using a novel first-class constructor context primitive. The functional and imperative algorithms are equivalent: we mechanise the proofs establishing this in the Coq proof assistant using the Iris framework. This yields a first fully verified implementation of well known algorithms on binary search trees with performance on par with the fastest implementations in C.

CCS Concepts: • **Software and its engineering** → **Control structures**; **Recursion**; • **Theory of computation** → **Operational semantics**.

Additional Key Words and Phrases: FIP, FBIP, Tail Recursion Modulo Cons, Zippers, Splay Trees, Zip Trees

ACM Reference Format:

Anton Lorenzen, Daan Leijen, Wouter Swierstra, and Sam Lindley. 2024. The Functional Essence of Imperative Binary Search Trees. *Proc. ACM Program. Lang.* 8, PLDI, Article 168 (June 2024), 25 pages. <https://doi.org/10.1145/3656398>

1 INTRODUCTION

In his book on purely functional data structures, Okasaki [1999b] presents several implementations of binary search trees. The inductive nature of these purely functional algorithms makes them amenable to reasoning and verification. A typical exercise in verification courses asks for a formal proof that insertions or deletions preserve the binary search tree properties [Appel 2018]. It is even possible to synthesize such implementations automatically [Albarghouthi et al. 2013; Polikarpova et al. 2016] or mechanically compute their amortized time complexity [Leutgeb et al. 2022; Schoenmakers 1993]. Unfortunately, the absolute performance of these functional algorithms is often worse than the imperative implementations published in papers and textbooks. While the functional algorithms have the same *semantic* behaviour as their imperative counterparts, their *operational* behaviour differs substantially. The functional algorithms suffer from asymptotically worse heap allocation (due to copying immutable data) and stack usage (due to non-tail calls).

Authors' addresses: Anton Lorenzen, University of Edinburgh, Edinburgh, UK, anton.lorenzen@ed.ac.uk; Daan Leijen, Microsoft Research, Redmond, WA, USA, daan@microsoft.com; Wouter Swierstra, Universiteit Utrecht, Utrecht, Netherlands, w.s.swierstra@uu.nl; Sam Lindley, University of Edinburgh, Edinburgh, UK, sam.lindley@ed.ac.uk.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2024 Copyright held by the owner/author(s).


ACM 2475-1421/2024/6-ART168

<https://doi.org/10.1145/3656398>

This paper presents novel functional algorithms on binary search trees with absolute performance on par with the best known imperative implementations. Enabled by recent advancements, including fully in-place functional programming [Lorenzen et al. 2023b] and Perceus reference counting [Reinking, Xie et al. 2021; Ullrich and de Moura 2019], our algorithms perform in-place updates whenever possible, without sacrificing their purely functional nature. Additionally, our algorithms execute in constant stack space: they are defined in a tail recursive manner, storing incomplete trees as heap-allocated *one-hole contexts*. Each of these algorithms is proven to be equivalent to the established functional version using a simple inductive argument. Yet these algorithms have a close operational correspondence to the published imperative code—we show how the one-hole contexts are an essential part of the loop invariant necessary to verify their imperative counterparts.

A one-hole context is typically represented by a zipper [Huet 1997], storing the path from the hole back up to the root. Our zipper-based algorithms navigate through the tree, accumulating a zipper of unvisited subtrees along the way. Upon encountering the key we were looking for, the zipper is unrolled to reconstruct a complete binary search tree. For example, our implementation of the ‘zig-zag’ case for splay trees (Section 5) looks something like this:

```
fip fun splay( accz, b, x, c )
  match accz
  NodeR(a,y,NodeL(up,z,d)) ->
    splay( up, Node(a,y,b), x, Node(c,z,d) )
  ...
```



The zipper we have accumulated, `accz`, stores the path back to the root, where each constructor `NodeR/NodeL` records whether the traversal down the tree went right or left. Our algorithm is readily calculated from the standard definition by means of a defunctionalized CPS-transformation [Danvy and Nielsen 2001; Reynolds 1972]. At the same time, it corresponds exactly to Figure 3 of Sleator and Tarjan [1985], given on the right. It is the essence of an imperative bottom-up algorithm, implemented with pointer reversal [Schorr and Waite 1967].

However, zippers are not the *only* way to represent one hole contexts. To give a functional account of imperative top-down algorithms, zippers are an asymptotically worse choice of data structure—much in the same way that lists are a poor implementation of queues. To address this, we introduce a novel language feature, *first-class constructor contexts*, that safely encapsulate the required mutation behind a purely functional interface. These contexts follow a design proposed by Minamide [1998], but we are the first to give an implementation that requires neither a linear type system nor reference counting. Using these constructor contexts, we can implement the corresponding zig-zag step of top down splay trees as follows:

```
fip fun splay( t, k, accl, accr )
  match t
  Node(ayzb,x,c) -> if x > k then match ayzb
  Node(a,y,zb) -> if y < k then
    splay(zb, k, accl ++ ctx Node(a,y,_),
          accr ++ ctx Node(_,x,c))
  ...
```



Here we traverse the tree and accumulate two constructor contexts `accl` and `accr`. Once again, this code can be calculated from the direct recursive definition, in this case using a tail recursion modulo context transformation [Leijen and Lorenzen 2023]. At the same time, it corresponds exactly to Figure 11 of Sleator and Tarjan [1985], given on the right, describing the imperative top-down algorithm that accesses a given key, whilst restructuring the tree in-place.

We illustrate our techniques by studying three particular implementations of restructuring binary search trees: *move-to-root* trees [Allen and Munro 1978; Stephenson 1980] (Section 2), self-adjusting *splay* trees [Sleator and Tarjan 1985] (Section 5) and the more recently published randomized *zip*

trees [Tarjan et al. 2021] (Section 6). For each of these data structures, we can derive both functional top-down and bottom-up algorithms whose performance is on par with the implementations proposed in the original papers. Yet correspondence of the functional and imperative versions goes even further: the functional algorithms capture the key loop invariants required to *verify* the imperative algorithms.

To make this precise, we formalize the imperative algorithms from the original papers in (a variant of) HeapLang, the imperative language used by the separation logic framework Iris [Jung et al. 2018]. After defining the loop-invariants relating the functional and imperative algorithms, functional correctness follows fully automatically thanks to our proof automation built with Diaframe [Mulder et al. 2023 2022], yielding a *novel and formal correctness proof* of these imperative algorithms. The proofs built on these loop invariants not only verify the correctness of the imperative algorithms, but also provide further evidence that the techniques presented here elucidate the inductive nature lurking hidden beneath their imperative shell: the functional essence of the imperative algorithms.

More specifically, we make the following contributions in this paper:

- We show that first-class constructor contexts (Section 3.2) enable true functional top-down algorithms to be implemented efficiently. We are the first to present an efficient implementation that does not require a linear type system or reference counting.
- We present several novel functional algorithms for binary search tree insertion (Sections 2, 5 and 6) and provide benchmarks that show that their performance in the Koka language [Leijen 2021] is on par with the best known imperative implementations written in C and several times faster than the corresponding implementations in OCaml and Haskell (Section 7).
- To the best of our knowledge, we give the first formal machine-checked correctness proofs for the imperative insertion algorithms of move-to-root, splay, and zip trees, exactly as published in the original papers (Sections 4, 5, and 6 together with Appendix C in the tech report).

Moreover, this work also gives novel insights into the design of imperative algorithms:

- We show that bottom-up algorithms can be calculated from a recursive specification using a defunctionalized CPS-transformation, while top-down algorithms can be derived using tail recursion modulo cons with product contexts (Section 2).
- We give new insights into splay trees by showing that they differ from move-to-root trees in only one balancing step, and characterizing how the bottom-up and top-down splay tree algorithms differ (Section 5).
- We derive an algorithm for top-down zip tree insertion which is simpler, but as efficient as the original algorithm given by Tarjan et al. [2021]. Furthermore, we also derive a bottom-up zip tree insertion algorithm. We are not aware of any previous bottom-up algorithm for zip tree insertion prior to the implementation given here (Section 6).

We include all programs and proofs in our artifact [Lorenzen et al. 2024].

2 MOVE-TO-ROOT TREES

To introduce our techniques, we consider *Move-to-root* trees, independently described by both Allen and Munro [1978] and Stephenson [1980], which are binary search trees, where accessing a particular key ensures it moves to the tree’s root. These trees are rarely used in practice since they can become unbalanced, but their simplicity makes them well suited to illustrate our key ideas.

2.1 A Recursive Functional Algorithm

All our examples in this paper are written in the Koka language [Leijen 2014] (v3.1.1). In Koka, we can declare a datatype for binary trees as:

```

type tree
  Node( left : tree, key : key, right : tree )
  Leaf

```

We use an abstract type `key` for the keys stored in the tree but this is usually instantiated to be an `int`. The main operation on binary trees is the `insert` function that takes a tree and a key as its arguments. If the key is not yet in the tree, the `insert` function inserts it; otherwise no elements are inserted or deleted. We can elegantly express move-to-root tree insertion inductively in a pure functional way using direct recursion, where we ensure in each branch that the inserted key always ends up at the root of the resulting tree:

```

fun insert( t : tree, k : key ) : tree
  match t
    Node(l,x,r) -> if x < k then match insert(r,k)
                        Node(s,y,b) -> Node( Node(l,x,s), y, b)
                        elif x > k then match insert(l,k)
                        Node(s,y,b) -> Node( s, y, Node(b,x,r))
                        else Node(l,k,r)
    Leaf -> Node(Leaf,k,Leaf)

```

This version performs a single traversal over the input tree. It is straightforward to verify various correctness properties of this function using a proof assistant such as Coq [2017]. For instance, we have proven that (1) whenever `t` is a binary search tree, so is `insert(t,k)`; (2) every key in `t` also occurs in `insert(t,k)`; and (3) the key stored at the root of `insert(t,k)` is equal to `k`.

Even though this recursive functional definition is simple enough, it does have its drawbacks. Firstly, it is not tail-recursive and can use stack space linear in the size of the tree in the worst-case. Depending on the implementation, its execution may also allocate many nodes in the process leading to (much) worse performance when compared to the imperative algorithms. We now proceed to derive efficient *fully in-place* bottom-up and top-down variants that remedy both these issues.

2.2 Bottom-Up Move-To-Root Using Zippers

A bottom-up algorithm first traverses down the tree to the insertion point and then restructures the tree on the way back up. This can be derived by a standard defunctionalized CPS transformation of the direct recursive definition of the `insert` function [Danvy and Nielsen 2001; Reynolds 1972]. Doing so, uncovers a familiar functional data structure, the *zipper* (or ‘one hole context’) [Huet 1997; McBride 2001] on binary trees:

```

type zipper
  Done
  NodeL(up : zipper, key : key, right : tree )
  NodeR(left : tree, key : key, up : zipper )

```

Note that the zipper stores the traversed path down *in reverse*, such that we can rebuild a new tree bottom-up in constant space in a single traversal:

```

fip fun rebuild( z : zipper, t : tree )
  match z
    Done -> t
    NodeR(l,x,up) -> match t // we came from the right
      Node(s,y,b) -> rebuild( up, Node( Node(l,x,s), y, b))
    NodeL(up,x,r) -> match t // we came from the left
      Node(s,y,b) -> rebuild( up, Node( s, y, Node(b,x,r)))

```

The `rebuild` function repeatedly moves up through the zipper, reassembling the original tree. This function may be executed *fully in-place*, as indicated by the `fip` keyword, a point we discuss more extensively in the next section. Using these definitions, we now specify the complete tail-recursive bottom-up `insert` function:

```

fig(1) fun insert-bu( t : tree, k : key )
  down-bu(t,k,Done)
fig(1) fun down-bu( t : tree, k : key, z : zipper )
  match t
  Node(l,x,r) -> if x < k then down-bu( r, k, NodeR(l,x,z) )
                elif x > k then down-bu( l, k, NodeL(z,x,r) )
                else rebuild( z, Node(l,k,r) )
  Leaf        -> rebuild( z, Node(Leaf,k,Leaf) )

```

The `insert-bu` function is both *fast* and *correct*. It is tail recursive; the `fig` property ensures no unnecessary new memory is allocated. In this case, the `fig(1)` annotation allows allocation of a single constructor when the inserted key is not yet present; otherwise no memory is allocated or freed. This addresses the two performance issues associated with the direct recursive definition we saw previously. To prove the `insert-bu` function is correct, we prove the following theorem relating the two algorithms:

Theorem 1. (*Correctness of bottom-up move-to-root insertion*)

The recursive- and the bottom-up algorithms coincide:

$$\text{down-bu}(t,k,z) \equiv \text{rebuild}(z, \text{insert}(t,k))$$

Proof. The proof proceeds by induction on the tree t . The base case, when t is a leaf, is trivial. If the tree is non-empty, we distinguish three cases, depending on the key x stored at t is less than, greater than, or equal to k . We cover the first case – the others are similar – where we need to show:

$$\text{down-bu}(r, k, \text{NodeR}(l,x,z)) \equiv \text{rebuild}(\text{NodeR}(l,x,z), \text{insert}(r,k))$$

which follows immediately from our induction hypothesis.

An obvious corollary of this theorem is that the recursive version calculated at the beginning of this section coincides with the tail-recursive bottom-up insert function presented here: for all trees t and keys k , we have $\text{insert-bu}(t,k) \equiv \text{insert}(t,k)$.

2.3 Intermezzo: First-Class Constructor Contexts

A top-down algorithm traverses a structure down in a single pass and directly returns the result structure when reaching the final position. Unfortunately, in a purely functional language is not always possible to express such algorithms directly. Consider the `map` function for example:

```

fun map( xs : list<a>, f : a -> b ) : list<b>
  match xs
  Cons(x,xx) -> Cons( f(x), map(xx,f) )
  Nil        -> Nil

```

Naively, this function would use stack space linear in the size of the first list. A well-known solution to derive a tail-recursive version is to use an *accumulator* for the result list, as:

```

fun map-acc( xs : list<a>, f : a -> b, acc : list<b> ) : list<b>
  match xs
  Cons(x,xx) -> map-acc( xx, f, Cons(f(x), acc) )
  Nil        -> reverse(acc)

```

$$\text{fun map}(xs, f) \quad \text{map-acc}(xs, f, \text{Nil})$$

Yet the resulting algorithm is no longer top-down, as eventually we need to traverse back through the accumulator to reverse it in $O(n)$ time. Similarly, using a function (or difference list [Clark and Tärnlund 1977; Hughes 1986]) as the accumulator, requires a final application of the functional accumulator, essentially traversing back up through the composite `Cons` operations.


To express true top-down algorithms, we introduce the concept of *first-class constructor contexts*. This abstraction can safely encapsulate the limited form of mutation necessary to define top-down algorithms, while still having a purely functional interface. We define a *constructor context* as a sequence of constructor applications that ends in a single *hole*. We can describe such contexts using the following grammar:

$$v ::= \dots \mid \text{ctx } K \quad K ::= \square \mid C^k v_1 \dots v_k$$

where we use v for values, and C^k for a constructor that takes k arguments. In Koka, the keyword `ctx` starts a constructor context and the hole \square is written as an underscore `_`. For example, we can write a list constructor context as `ctx Cons(1,_)` or a binary tree constructor context as `ctx Node(Node(Leaf,1,Leaf),2,_)`.

Constructor contexts support two operations: we can compose (or “append”) two contexts, written $c_1 \# c_2$, or apply a value to a context to fill the hole, written $(c \# v)$. For example, the expression `(ctx Cons(1,_) ++ (ctx Cons(2,_) ++ Nil)` evaluates to `(ctx Cons(1,Cons(2,_) ++ Nil)` and then to `[1,2]`. Similarly:

`(ctx Node(Node(Leaf,1,Leaf),2,_) ++ (ctx Node(_,4,Node(Leaf,5,Leaf)))) ++ Node(Leaf,3,Leaf)` appends and applies binary tree contexts. This can be visualised as follows:



A natural implementation of contexts is as lambda expressions where context composition and application correspond to function composition and application:

$$\text{ctx } K = \lambda x. K[x] \quad \text{with } x \notin \text{fv}(K) \qquad c_1 \# c_2 = c_1 \circ c_2 \qquad c \# e = c e$$

We will sometimes use this naive implementation when *reasoning* about programs, but it is rather inefficient. Section 3.2 presents a novel technique for implementing contexts efficiently where composition and application become constant time operations.

Using constructor contexts, we are now able to define a true top-down functional version of `map` using an *accumulating constructor context*:

```
fun map-td( xs : list<a>, f : a -> b, acc : ctx<list<b>> ) : list<b> fun map(xs,f)
  match xs
  Cons(x,xx) -> map-td( xx, f, acc ++ ctx Cons(f(x),_) )      map-td(xs, f, ctx _)
  Nil        -> acc ++ Nil
```

The `map` function now uses a single tail-recursive traversal down the list and returns the final list directly (in constant time) when reaching the end of the list.

2.4 Top-Down Move-To-Root Using Accumulating Constructor Contexts

Using accumulating constructor contexts, we now also define a true top-down functional version of move-to-root insertion. The very first recursive version of `insert` matches on the result of a recursive call, adding a constructor to either the left or right subtree:

```
... match insert(r,k)
  Node(s,y,b) -> Node( Node(l,x,s), y, b)
```

We can make this tail-recursive by passing *two* accumulating constructor contexts for each of the left- and right subtrees of the root. For example, in the case outlined above, we extend the context we have accumulated so far with an additional `ctx Node(1,x,_)`. In the base cases, we then use the accumulated contexts to construct the final tree:

```
fip(1) fun insert-td( t : tree, k : key ) : tree
  down-td(t,k,ctx _, ctx _)

fip(1) fun down-td( t : tree, k : key, accl : ctx<tree>, accr : ctx<tree> ) : tree
  match t
  Node(l,x,r) -> if x < k then down-td( r, k, accl ++ ctx Node(1,x,_) , accr )
                elif x > k then down-td( l, k, accl, accr ++ ctx Node(_,x,r) )
                else Node( accl ++ l, x, accr ++ r )
  Leaf       -> Node( accl ++ Leaf, k, accr ++ Leaf )
```

We start with two empty contexts `ctx _`, and we can see that in each branch we extend either the left context with a tree of smaller elements, or the right context with a tree of larger elements. There is

no need to traverse a zipper “back up” as we did for the bottom-up algorithm. It is straightforward to prove the following equality, relating the recursive version of move-to-root trees calculated at the beginning of this section to top-down version using constructor contexts:

Theorem 2. (*Correctness of top down move to root insert*)

For all trees t , keys k and constructor contexts $accl$ and $accr$,

$$\text{down-td}(t, k, accl, accr) \equiv \text{val Node}(l, x, r) = \text{insert}(t, k) \text{ in Node}(accl ++ l, x, accr ++ r)$$

Proof. The proof proceeds by induction on the tree t . The base case, when t is a leaf, is trivial. If the tree is non-empty, we distinguish three cases, depending on the key x stored at t is less than, greater than, or equal to k . For the first case, we need to show:

$$\begin{aligned} \text{down-td}(r, k, accl ++ \text{Node}(l, x, _), accr) &\equiv \\ \text{val Node}(l', x', r') = \text{insert}(r, k) \text{ in Node}(accl ++ \text{Node}(l, x, _) ++ l', x', accr ++ r') &\equiv \end{aligned}$$

This follows from our induction hypothesis for the right subtree, where we use the (*dist*) law. The other cases are similar. A corollary of this result is that that the recursive insertion coincides with top-down insertion algorithm: for all trees t , keys k , $\text{insert-td}(t, k) \equiv \text{insert}(t, k)$

3 MAKING IT FAST: FULLY IN-PLACE AND CONSTRUCTOR CONTEXTS

The previous section has shown three different implementations of the same algorithm: a direct recursive definition; an efficient bottom-up implementation using a zipper; and an efficient top-down implementation using first-class constructor contexts. To achieve performance that is competitive with the imperative algorithms, we need two fundamental techniques: fully in-place functions (*flip*) and efficient constructor context operations.

3.1 Fully In-Place Functional Programming

Throughout this paper, we write purely functional programs, but the goal is always to derive *fully in-place* or *flip* functions that can be compiled to efficient code. This section highlights the key principles underlying this recent paradigm of fully-in place functional programming [Lorenzen et al. 2023b]. Consider for example the function that swaps the left and right subtrees:

```
flip fun swap (t : tree) : tree
  match t
  | Leaf          -> Leaf
  | Node(l, x, r) -> Node(r, x, l)
```

Lorenzen et al. [2023b] define a linear calculus characterizing *flip* programs and prove that any program in the *flip* fragment can be compiled to code that does not use any (de)allocations and uses constant stack space: it can be executed fully in-place. The *flip* keyword asserts that a given function is in this linear fragment. The Koka compiler statically checks that each *flip* function does not duplicate or discard its arguments; when a function is erroneously marked as *flip*, the Koka compiler gives a warning statically.

Intuitively, we can check if a function is *flip* by ensuring that in each branch the constructors matched on use the same memory layout as the constructors “allocated” on the right hand side of the function definition, thereby ensuring every heap cell is reused. The reuse analysis allows constructors from different datatypes to reuse the same memory cells, illustrated by the following case from the earlier *down-bu* function:

```
Node(l, x, r) -> if x < k then down-bu( r, k, NodeR(l, x, z))
```

More formally, in each branch of a case expression, the constructor that is matched provides us with a *reuse credit* of a certain size k , written as \diamond_k (similar to the space credits of Aspinall et al. [2008]). These reuse credits are consumed when space of that size is required: in *down-bu* the *NodeR*(l, x, z) consumes the reuse credit \diamond_3 obtained by matching on the *Node* constructor. Constructors without arguments, like *Nil*, *True*, or *Leaf*, and primitive types like integers, are called *atoms* and require no

allocation. Furthermore, value types like tuples are always unboxed and passed as registers or on the stack.

Nevertheless, it is only safe to reuse these memory locations if the original parameters are *owned* and unique at runtime! Inside `fib` functions the linear use of owned parameters is guaranteed, but when `fib` functions are called from a non-`fib` context, the arguments may be shared. Consider the following example:

```
fun mirror( t : tree, k : key ) : tree
  Node(t,k,swap(t))
```

Here the tree `t` is now shared. Any in-place update on `t` would be unsound and change the meaning of this program. To ensure `fib` functions are executed safely, Koka uses precise reference counting [Reinking, Xie et al. 2021; Ullrich and de Moura 2019] to determine dynamically whether or not arguments can be reused in-place. In particular, for a function like `swap`, the generated code becomes:

```
fib fun swap(t : tree) : tree
  match t
  Leaf          -> Leaf
  Node(l,x,r) -> val p = if unique(t) then &t else {dup(l); dup(r); decref(t); alloc(3)}
                in Node@p(r,x,l)
```

That is, if `t` does have a unique reference count of 1 at runtime, the allocated space is reused. Otherwise, `t` is shared: the reference counts are adjusted and a fresh heap cell is allocated.

The `fib` annotation in Koka only guarantees that no (de)allocation occurs if the parameters are unique at runtime. This may be viewed as weakness – we do not guarantee statically that a function will actually be executed in-place – but it does offer greater flexibility where we can use `fib` functions in both modes. In particular, for the tree algorithms in this paper, we not only get the efficient in-place updating behaviour for unique trees, but we can also use them *persistently* where any shared (sub)trees are copied as needed.

The `fib` check provides a strong guarantee: constant stack usage and no (de)allocation at all. Throughout this paper, we also use the `fib(n)` variant which allows a function to allocate at most `n` constructors. This is useful for tree insertion algorithms, as we may need to allocate a constant amount of memory for the single node storing the new key.

Improving Bottom-Up. The `swap` function may seem trivial – but consider the following slight variation that rotates a binary tree, moving subtrees from the left to the right:

```
fib fun rotate-right( t : tree ) : tree
  match t
  Node(Node(l1,lx,lr),x,r) -> Node(l1,lx,Node(lr,x,r))
  Node(Leaf,x,Leaf)       -> Node(Leaf,x,Leaf)
  Leaf                    -> Leaf
```

It is easy to check that this function is fully in-place. As `fib` functions can safely call other `fib` functions, we can rewrite our `rebuild` function as follows:

```
fib fun rebuild( z : zipper, t : tree ) : tree
  match z
  Done -> t
  NodeL(up,x,r) -> rebuild( up, rotate-right(Node(t,x,r)) )
  NodeR(l,x,up) -> rebuild( up, rotate-left(Node(l,x,t)) )
```

This now corresponds closely to the published algorithm by Allen and Munro [1978] where they also use a bottom-up traversal using left- and right rotations. We formalise the precise relation between our bottom-up `insert-bu` function and the published algorithms shortly (Section 4), but first turn our attention to the top-down version of the same algorithm.

3.2 Efficient First Class Constructor Contexts

As mentioned previously, constructor contexts can be implemented using functions, but such implementation is unnecessarily inefficient. Minamide [1998] describes a linear hole calculus for constructor contexts. In Minamide's system, a context has an affine type and it is safe to update the hole *in-place*. A context is represented by a *Minamide tuple*, written as $\{x, h\}$, where the first element x points to the top of the data structure, and the second element h points directly to the hole inside that structure. Composition and application can now directly update the hole in-place and are constant time.

Unfortunately, it is not easy to extend an existing language with Minamide's system as it requires an affine type system for contexts (and also uses linear derivations and evaluation under-lambda for contexts). In particular, this is problematic for some of the proofs we do in this paper that rely on referential transparency and do not rule out sharing or duplication of contexts.

Context Paths. There is a way though to have efficient in-place mutating context operations *without* requiring affine types. The key to this is the use of runtime *context paths*, which store the path from the root to the hole, first described by Leijen and Lorenzen [2023].

Their use of context paths is internal and not exposed to the user, but we can use a similar runtime mechanism to implement our first-class constructor contexts.

In essence, we compile constructor contexts to a runtime representation storing the context path down from the top to the hole in the data structure. To enable this, we use extra bits in the header of each object where we store the index of the child that leads to the (single) hole in the structure. Koka re-uses an 8-bit field for this purpose which is normally used for stackless freeing.

The context path indices can be constructed in constant time when compiling constructor contexts. Writing C_i for the constructor C decorated with child index i , we compile a constructor context into a Minamide tuple as follows:

$$\begin{aligned} \text{ctx } \square &= \{\square, \text{NULL}\} \\ \text{ctx } C \ x_1 \dots x_{i-1} \ \square \ x_{i+1} \dots x_k &= \text{let } x = C_i \ x_1 \dots x_{i-1} \ \square \ x_{i+1} \dots x_k \ \text{in } \{x, \&x.i\} \\ \text{ctx } C \ x_1 \dots x_{i-1} \ K \ x_{i+1} \dots x_k &= \text{let } \{x, h\} = \text{ctx } K \ \text{in } \{C_i \ x_1 \dots x_{i-1} \ x \ x_{i+1} \dots x_k, h\} \quad (K \neq \square) \end{aligned}$$

where $\&x.i$ denotes the address of field i in x . At runtime, a constructor allocation of C typically initializes the header fields, including the tag. Adding in the context path index yields a single constant, eliminating any overhead associated with this representation. For example, the Koka compiler compiles a context like `ctx Node(Node(Node(Leaf, 1, Leaf), 2, _), 5, Leaf)` internally into:

```
val x = Node3(Node(Leaf, 1, Leaf), 2, □) in { Node1(x, 5, Leaf), &x.3 }
```

where each constructor along the context path is annotated with a child index (1 and 3).

With these context paths, we can now follow the path from the top of a context to the hole in that structure *at runtime*, and thus we are able to *copy* the linear context path dynamically at runtime when required. When we compose or apply a context we can now copy shared contexts only when needed. In a language with precise reference counts (like Koka or Lean) we copy the contexts at runtime along the context paths whenever they are not unique.

We can also support this in languages without precise reference counts though. In particular, we can use a special distinguished value for a runtime hole \square that is never used by any other object. A substitution now first checks the value at the hole: if it is a \square value, the hole is substituted for the first time and we just overwrite the hole in-place (in constant time). However, any subsequent substitution on the same context will find some object instead of \square . At this point, we first dynamically copy the context path (in linear time) and then update the copy in-place.

If the contexts happen to be used linearly, then all operations execute in constant time, just as in Minamide's approach; but we now have full functional semantics and any subsequent substitutions on the same context work correctly (but will take linear time in the length of the context path). The




Fig. 1. Applying a shared context where the context path is denoted with a bold edge. In the second update the nodes along the context path are copied.

expression `val c = ctx Cons(1,_) in (c ++. [2] , c ++. [3])`, where the context c is shared, evaluates correctly to $([1, 2], [1, 3])$.

Figure 1 illustrates a more complex example of a shared tree context that is applied to two separate nodes. In the illustration the runtime context path is denoted by bold edges. The intermediate state is interesting as it is both a valid tree, but also a part of the tree is shared with the remaining context, where the hole points to a regular node now. When that context is applied, only the context path (node 5 and 2) is copied first where all other nodes stay shared (in this case, only node 1).

However, in the context composition operation $c_1 \# c_2$ we need an extra check in order to avoid cycles: we check if c_2 has an already overwritten hole or if the hole in c_2 is at the same address as in c_1 . In either case, c_2 is copied along the context path. Effectively, both checks ensure that the new context that is returned always ends with a single fresh `HOLE`. If we compose a context with itself:

```
val c = ctx Cons(1,_) in (c ++ c) ++. [2]
```

this evaluates to $[1, 1, 2]$, where the check copies the appended c . In Appendix D in the tech report, we give an implementation of these constructor contexts in C and prove that these checks are sufficient for avoiding cycles using a heap semantics.

4 FUNCTIONAL AND IMPERATIVE MOVE-TO-ROOT COINCIDE

In Section 2, we derived two functional implementations from our recursive specification. How can we relate these to the imperative move-to-root algorithms published by Stephenson [1980] and Allen and Munro [1978]? As we will see, the imperative algorithms rely heavily on pointer manipulation: it is not at all obvious that they are correct or even represent the ‘same’ program.

These published algorithms are usually written in imperative pseudocode. To reason about them, we formalize each algorithm precisely in Iris [Jung et al. 2018], a framework for (higher-order concurrent) separation logic [Reynolds 2002] implemented as a library in the Coq proof assistant [2017]. In the style of Frumin et al. [2019] and Bedrock2 [Erbsen et al. 2021; Pit-Claudel et al. 2022], we have defined an embedded language, called *AddressC*, building on the standard *HeapLang* language supported by Iris [2022].

The *AddressC* language is embedded into Coq where we use extensive *Notation* to have the embedded code resemble a low-level C-style language that can match the typical pseudo-code in published algorithms closely. Eventually, *AddressC* is desugared into a standard *HeapLang* value representing the low-level control-flow and heap operations on which the proofs operate.

While our language builds on *HeapLang*, we place special consideration on precisely modeling while loops and the (untyped) low-level structure of memory. For example, we model a tree as:

```
Fixpoint is_tree (t : tree) (v : val) : iProp Σ :=
  match t with
  | Leaf => 「v = NULL」
  | Node l x r => ∃(p:loc) l' r', 「v = #p」 * p ↦* [ l'; #x; r' ] * is_tree l l' * is_tree r r'
  end.
```

```

Definition heap_mtr_insert_td : val :=
fun ( name, root ) {
  var: left_dummy := NULL in
  var: right_dummy := NULL in
  var: node := root in
  var: left_hook := &left_dummy in
  var: right_hook := &right_dummy in
  while: ( true ) {
    if: ( node != NULL ) {
      if: ( node->value == name ) {
        *left_hook = node->left;;
        *right_hook = node->right;;
        root = node;;
        break
      }
      else {
        if: ( node->value > name )
        {
          *right_hook = node;;
          right_hook = &(node->left);;
          node = node->left
        }
        else
        {
          *left_hook = node;;
          left_hook = &(node->right);;
          node = node->right
        }
      }
    }
    else {
      *left_hook = NULL;;
      *right_hook = NULL;;
      root = AllocN #3 NULL;;
      root->value = name;;
      break
    }
  };
  root->left = left_dummy;;
  root->right = right_dummy;;
  ret: root
}.

```

```

node := root;
left_hook := addr(left(dummy));
right_hook := addr(right(dummy));

while node ≠ null do
  if value(node) = name then
    begin
      0(left_hook) := left(node);
      0(right_hook) := right(node);
      root := node;
      go to bottom
    end;
  if value(node) > name then
    begin
      0(right_hook) := node;
      right_hook := addr(left(node));
      node := left(node)
    end
  else
    begin
      0(left_hook) := node;
      left_hook := addr(right(node));
      node := right(node)
    end;
end;

0(left_hook) := null;
0(right_hook) := null;
root := new_node ( );
value(root) := name;

bottom:
left(root) := left(dummy);
right(root) := right(dummy)

```

Fig. 2. The move-to-root top-down algorithm formalized in AddressC on the left, versus a screenshot of Stephenson’s published algorithm on the right (© Springer Nature)

This states that a *Leaf* is represented by a null address (the constant `NULL`). To represent a non-empty tree, `Node(1, x, r)`, requires having some address `#p`, pointing to a heap cell of 3 fields containing an address for the left tree (`1'`), its key `#x`, and an address for the right tree `r'`. The separating conjunction, `*`, ensures that the tree is indeed inductively defined and that there are no cycles. For the bottom-up algorithms we additionally need to model zippers, which requires us to distinguish between `NodeL` and `NodeR`. To do so, we include an additional tag field in the heap cells, as $p \mapsto * [\#1; 1'; \#x; r']$.

Variables typically denote memory locations, and just as in C, we use `&` to take an address of a location and we use `*` to dereference an address. We can also dereference at an offset, writing `node[[2]]` to dereference the second field of the `node` address. We usually define notation for constant offsets so we can write `node->right` instead of `node[[2]]` to get the value of the right child.

4.1 Proving Stephenson’s Top-Down Algorithm

Stephenson [1980] presents an imperative top-down insertion algorithm for top-down move-to-root trees in pseudocode. Figure 2 shows both Stephenson’s top-down algorithm as published and our formal AddressC implementation. We can see that our formal AddressC implementation corresponds to the published algorithm almost line-by-line. Using Iris, we can now formally relate the functional algorithm and AddressC implementation:

Theorem 3. (Stephenson’s imperative top-down move-to-root algorithm is correct)

```

Lemma heap_mtr_insert_td_correct (k : Z) (tv : val) (t : tree) :
  {{{ is_tree t tv }}}
  heap_mtr_insert_td (ref #k) (ref tv)
  {{{ v, RET v; is_tree (mtr_insert_td t k) v }}}.

```

The pre-condition requires that the argument address tv points to a valid in-memory tree corresponding to t , and the post-condition establishes that the result address v points to a valid in-memory tree that corresponds to $mtr_insert_td\ t\ k$. The proof goes through because we can directly relate the loop invariant of this algorithm to the recursive calls of mtr_insert_td . As we see in the next Section, this is only possible because constructor contexts *precisely* capture the top-down behaviour of Stephenson’s algorithm. It would be much harder to relate the AddressC code to our original recursive definition. As is often the case in verification, finding the right formulation of our theorem is vital – this proof would not be possible without constructor contexts.

4.2 Representing Constructor Contexts

Stephenson’s algorithm uses intricate pointer manipulation and even goto-statements that make it non-trivial to verify formally. The key insight is that Stephenson builds the *smaller* and *bigger* trees using the `left_hook` and `right_hook` variables. For example, for the case where the key in the current node is larger than the argument key (`name`), we have:

```
if: ( node->value > name ) {
  *right_hook = node;;
  right_hook = &(node->left);;
  node = node->left }
```

Here the current node address is written to `*right_hook` which is then itself updated to point to the left child of the current node (`right_hook = &(node->left)`). Afterwards the current node is advanced to the left child. This corresponds to the functional version, where the current node is written into the right context (`accr`) and the hole is set to its left child:

```
down( l, k, accl, accr ++ ctx Node(_,x,r) )
```

At this point though, we have all kinds of problems in the formal setting. Not only have we overwritten the value that `right_hook` was previously pointing to, but we have introduced aliasing where both the current *bigger* tree’s left-child and `node` point to the same location. The *bigger* tree is not even a valid constructor context as the left child is “dangling” pointer (that will eventually be overwritten). Yet we can still prove these pointer manipulations correct by relating them to the constructor contexts used in our functional algorithm.

First, we implement our functional algorithms in Coq using a slow, but purely functional representation of constructor contexts, where `append` and `composition` take time linear in the depth of the first context. Such a context can be modelled similar to a zipper, but with the pointers going from the root to the hole:

```
Inductive ctx : Set :=
| Node0 (l : ctx) (x : Z) (r : tree)
| Node2 (l : tree) (x : Z) (r : ctx)
| Hole.

Fixpoint comp (z1 : ctx) (z2 : ctx) : ctx :=
match z1 with
| Node0 z1 x r => Node0 (comp z1 z2) x r
| Node2 l x zr => Node2 l x (comp zr z2)
| Hole => z2 end.
```

We can then define an `is_ctx z p h` predicate that represents the context z in heap memory, with root pointer p and hole pointer h . However, we take care to ensure that the predicate does not take ownership of the hole h . This is different from the usual presentation [Charguéraud 2016] and allows us to change the value of the hole without inspecting the constructor context (to allow temporarily for a dangling pointer). For example, we can now prove the following lemma:

```
Lemma ctx_of_ctx (z1 : ctx) (z2 : ctx0) (zv1 : loc) (hv1 : loc) (zv2 : loc) (hv2 : loc) :
  is_ctx z1 zv1 hv1 * hv1  $\mapsto$  #zv2 * is_ctx0 z2 zv2 hv2  $\neg$  is_ctx (comp z1 z2) zv1 hv2.
```

This states that if the hole points to another context, together, they form the composed context. This is the key lemma that enables checking the individual cases of the algorithms.

Loop Invariant. Using our representation of constructor contexts, we can now show the full proof of the functional correctness of Stephenson’s top-down algorithm. We start the proof by

introducing the variables:

Proof.

```
wp_begin "Ht"; name, root. wp_var left_dummy. wp_var right_dummy.
wp_load. wp_var node. wp_var left_hook. wp_var right_hook.
```

The main difficulty of this proof lies in specifying the loop invariants for the while-loop. The first formula passed to the `wp_while_true` tactic gives the condition once the loop terminates and the second formula gives the invariant for subsequent iterations.

```
wp_while_true "H"
(∃ l (x : Z) r left_dummy_v right_dummy_v (p : loc) left_dangling_v right_dangling_v,
 root ↦ #p * p ↦* [left_dangling_v; #x; right_dangling_v]
 * left_dummy ↦ left_dummy_v * is_tree l left_dummy_v
 * right_dummy ↦ right_dummy_v * is_tree r right_dummy_v
 * 「mtr_insert_td t k = Node l x r」)%I
```

This first formula mirrors the *return value* of the functional code, which returns two trees `l` and `r`. It specifies that `left_dummy` and `right_dummy` point to those trees, that `root` points to an allocation containing `x` and that this final result `Node l x r` is equal to the result of the functional code `mtr_insert_td t k`. Our second formula is:

```
(∃ accl accr t' (left_hole right_hole : loc) left_hole_v right_hole_v root_v node_v,
 node ↦ node_v * is_tree t' node_v
 * root ↦ root_v * name ↦ #k
 * left_hook ↦ #left_hole * left_hole ↦ left_hole_v * is_ctx accl left_dummy left_hole
 * right_hook ↦ #right_hole * right_hole ↦ right_hole_v * is_ctx accr right_dummy right_hole
 * 「mtr_insert_td t k = mtr_down_td t' k accl accr」)%I.
```

The invariant for subsequent iterations mirrors the *recursive calls* of the functional code, which calls itself in tail-position on a tree `t'` and two contexts `accl`, `accr`. The invariant specifies that `left_dummy` and `right_dummy` point to these contexts and `left_hook` and `right_hook` point to the holes, while `node` points to the subtree and `name` stays constant. Finally, it asserts ownership over the root location and asserts that the functional values `accl`, `accr`, `t'` correspond to a loop iteration of the functional code `mtr_down_td`.

With these invariants, the remaining proof of the top-down algorithm is highly automated and we resolve all obligations associated with assignments on the heap or constructor contexts automatically using Diaframe’s `iSteps` tactic [Mulder et al. 2023 2022], which performs proof search guided by custom hints (such as our lemma `ctx_of_ctx`):

```
- iDecompose "H". iSteps. (* The claim follows once the while-loop ends *)
- iDecompose "H". rewrite H1. (* Rewrite the functional recursive call *)
  unfold array. repeat (iSteps; try case_bool_decide). (* Then the invariant is maintained *)
- iSteps. (* The invariant holds when entering the loop *)
```

Qed.

The brevity of the proof – despite the intricate nature of Stephenson’s algorithm – provides further evidence that these definitions capture the essence of top-down algorithms.

4.3 Proving Allen and Munro’s Bottom-Up Algorithm

While Allen and Munro [1978] do not present imperative pseudo-code, we can define an imperative version of their algorithm in AddressC. They introduce a “simple exchange” (corresponding to what is now called a rotation) and describe their algorithm as:

[...] perform a sequence of simple exchanges on the retrieved record so that it is moved to the root. ... By carefully using the coding trick of “reversing the direction of the pointers” in performing the search, only two or three extra storage locations are required.

We can directly implement the mentioned pointer reversal technique [Schorr and Waite 1967] in AddressC (see Appendix E.1 in the tech report). The code corresponds closely to the functional bottom-up version we derived earlier. In particular, just as constructor contexts captured the

top-down behavior in Stephenson’s algorithm, a *zipper* captures the structure of in-place pointer-reversal. Even though the use of pointer-reversal is complex from a formal perspective, we can use the functional zippers to relate the functional and imperative versions to make the proofs go through.

Theorem 4. (*Allen and Munro’s imperative bottom-up move-to-root algorithm is correct*)

```
Lemma heap_rebuild_correct (z : zipper) (t : root_tree) (zipper tree : loc) (zv tv : val) :
  {{{ zipper ↦ zv * is_zipper z zv * tree ↦ tv * is_root_tree t tv }}}
  heap_rebuild #zipper #tree
  {{{ v, RET v; is_root_tree (move_to_root z t) v }}}.
```

```
Lemma heap_mtr_insert_bu_correct (i : Z) (tv : val) (t : tree) :
  {{{ is_tree t tv }}}
  heap_mtr_insert_bu (ref #i) (ref tv)
  {{{ v, RET v; is_tree (mtr_insert_bu i t) v }}}.
```

The precondition of `heap_mtr_rebuild` requires that the argument addresses, `zipper` and `tree`, point to a zipper `z` and non-leaf binary tree `t`. The postcondition guarantees that after execution, the memory location that is returned, `v`, denotes the non-leaf binary tree arising from our functional algorithm, `rebuild`. Similarly, `heap_mtr_insert_bu_correct` specifies that given an arbitrary binary tree `t` with its heap representation `tv`, the imperative version returns a tree corresponding to `mtr_insert_bu i t`.

5 SPLAY TREES

Having looked at *move-to-root* trees, we can apply the same techniques to their improved sibling, *splay* trees [Sleator and Tarjan 1985]. The *move-to-root* trees only move the accessed element to the root of the tree but they do not restructure the tree. As such, the tree can degrade to a list in the case of ordered accesses. Splay trees on the other hand are self-adjusting: accessing an element also restructures the path to the root to become more balanced. Sleator and Tarjan [1985] identify six different kinds of tree rotations that are required, *zig*, *zigzig*, *zigzag* and their mirrored counterparts – is it possible to derive all of these rotations?

5.1 The Essence of Splay Tree Rebalancing

To understand the essence of splay trees, we look again at the definition of *move-to-root* trees since splay trees satisfy the exact same requirements. First, we unroll the recursion in *move-to-root* tree insertion once:

```
fun insert( t : tree, k : key )
  match t
  Node(l,x,r) ->
    if x < k then match r
      Node(r1,rx,rr) ->
        if rx < k then match insert(rr,k)
          Node(s,y,b) -> Node( Node(l,x,Node(r1,rx,s)), y, b) // (A)
        elif rx > k then match insert(r1,k)
          Node(s,y,b) -> Node( Node(l,x,s), y, Node(b,rx,rr))
      ...
```

Now we gain insight into why *move-to-root* trees can easily become unbalanced: when we move twice to the right (and dually, twice to the left for the $x > k$ case) as in the branch labelled (A), we create a short unbalanced part with two right leaning nodes:

```
(Node(l,x,Node(r1,rx,rr)) -> Node(l,x,Node(r1,rx,s) // A
```

A splay tree though *rotates* those nodes instead, resulting in a more balanced result:

```
(Node(l,x,Node(r1,rx,rr)) -> Node(Node(l,x,r1),rx,s)
```

This is the essence of splay tree restructuring – and it captures the key difference between *move-to-root* trees and splay trees. Using this insight, we obtain a recursive `insert` function (Appendix B in the tech report) that closely mirrors the version presented by Okasaki [1999b] (Sec. 5.4).

5.2 Top-Down Splay Trees

Just as with move-to-root trees (Section 2.4), we can again use accumulating constructor contexts to build the ‘smaller’ and ‘bigger’ trees on the way down. In particular, a match on a recursive call:

```
match insert(rr,k)
  Node(s,y,b) -> Node(Node(Node(l,x,r1),rx,s),y,b) // (A)
```

can be changed into a direct tail-recursive call, accumulating the constructor contexts:

```
splay(rr, k, accl ++ ctx Node(Node(l,x,r1),rx,_), accr)
```

The derived top-down version becomes:

```
fun insert-td( t : tree, k : key ) : tree
  down-td( t, k, ctx _, ctx _)
```

```
fip(1) fun down-td(t : tree, k : key, accl : ctx<tree>, accr : ctx<tree> ) : tree
  match t
    Node(l,x,r) ->
      if x < k then match r
        Node(r1,rx,rr) ->
          if rx < k then down-td(rr, k, accl ++ ctx Node(Node(l,x,r1),rx,_), accr)
          elif rx > k then down-td(r1, k, accl ++ ctx Node(l,x,_), accr ++ ctx Node(_,rx,rr))
          else Node( accl ++. Node(l,x,r1), rx, accr ++. rr )
        Leaf -> Node( accl ++. Node(l,x,Leaf), k, accr ++. Leaf )
      elif x > k then match l
        ...
      else Node( accl ++. l, x, accr ++. r )
    Leaf -> Node( accl ++. Leaf, k, accr ++. Leaf )
```

Now we have an efficient `fip(1)` function again. We can also formally check that top-down and direct splay tree insertion coincide:

Theorem 5. (Correctness of top-down splay tree insertion)

$\text{down-td}(t,k,\text{accl},\text{accr}) \equiv \text{val Node}(l,x,r) = \text{insert}(t,k) \text{ in Node}(\text{accl} ++. l, x, \text{accr} ++. r)$

5.3 Bottom-Up Splay Trees

We can also derive a bottom-up version by reusing the zipper datatype for move-to-root trees to keep track of the path we took when descending down the tree:

```
fip(1) fun insert-bu-fused( t : tree, k : key ) : tree
  down-bu-fused(t,k,Done)

fip(1) fun down-bu-fused( t : tree, k : key, z : zipper ) : tree
  match t
    Node(l,x,r) ->
      if x < k then down-bu-fused( r, k, NodeR(l,x,z) )
      elif x > k then down-bu-fused( l, k, NodeL(z,x,r) )
      else rebuild(z, Node(l,x,r))
    Leaf -> rebuild(z, Node(Leaf,k,Leaf))
```

All the restructuring now takes place in the `rebuild` function which matches two zipper nodes at a time to cover all rebalancing cases:

```
fip fun rebuild( z : zipper, t : root ) : tree
  match root
    Root(tl,tx,tr) -> match z
      Done -> Node(tl,tx,tr)
      NodeR(r1,rx,Done) -> Node(Node(r1,rx,t1),tx,tr) //zig
      NodeL(Done,lx,lr) -> Node(tl,tx,Node(tr,lx,lr))
      NodeR(r1,rx,NodeR(l,x,up)) -> rebuild( up, Root(Node(Node(l,x,r1),rx,t1),tx,tr) ) //zigzig
      NodeL(NodeR(l,x,up),lx,lr) -> rebuild( up, Root(Node(l,x,t1),tx,Node(tr,lx,lr)) )
      NodeR(r1,rx,NodeL(up,x,r)) -> rebuild( up, Root(Node(r1,rx,t1),tx,Node(tr,x,r)) ) //zigzag
      NodeL(NodeL(up,x,r),lx,lr) -> rebuild( up, Root(tl,tx,Node(tr,lx,Node(lr,x,r))) )
```

Even though the essence of splay tree restructuring has just two cases for rebalancing, the `rebuild` function now shows the usual six rebalancing cases that are common in the splay tree literature: *zig*, *zigzig*, *zigzag*, and their mirrored counterparts.

We call this the *fused* algorithm and it corresponds directly to Sleator and Tarjan’s published bottom-up algorithm. However, even though it splays correctly it turns out to be not equivalent to the recursive- or top-down algorithms. For example, if we start from a right-unbalanced tree with nodes 1 to 4 and insert node 4, we get different results for each of the various algorithms:

initial tree: move-to-root: top-down splay: (fused) bottom-up splay:



That the imperative top-down and bottom-up splay algorithms are not equivalent was not widely known [Lucas 2004], but when trying to prove equivalence on the functional specifications it becomes immediately apparent that $\text{down-bu-fused}(t, k, z) \equiv \text{rebuild}(z, \text{insert}(t, k))$ does not hold.

5.4 Functional and Imperative Splay Trees Coincide

Using our AddressC embedded language in Iris we formalized the published top-down and bottom-up algorithms by Sleator and Tarjan [1985], where we use pointer-reversal for bottom-up. Once again, there is a line-by-line correspondence between the published pseudocode and formal AddressC implementations that we have written (Appendix E.2 and E.3 in the tech report). Using the same techniques as for move-to-root trees we have formally established that the functional implementations accurately model the published imperative algorithms:

Theorem 6. (*Sleator and Tarjan’s imperative top-down splay algorithm is correct*)

```
Lemma heap_splay_insert_td_correct (k : Z) (tv : val) (t : tree) :
  {{{ is_tree t tv }}}
  heap_splay_insert_td (ref #k) (ref tv)
  {{{ v, RET v; is_tree (splay_insert_td k t) v }}}.
```

Theorem 7. (*Sleator and Tarjan’s imperative bottom-up splay algorithm is correct*)

```
Lemma heap_splay_insert_bu_correct (k : Z) (tv : val) (t : tree) :
  {{{ is_tree t tv }}}
  heap_splay_insert_bu (ref #k) (ref tv)
  {{{ v, RET v; is_tree (splay_insert_bu_fused k t) v }}}.
```

It is worth repeating that the proofs of these theorems are *direct*, requiring no additional lemmas. This is possible because the functional implementations precisely capture the iterative behaviour of their imperative counterparts through constructor contexts and zippers. Furthermore, these results are novel—to the best of our knowledge there is no formal correctness proof of these algorithms.

6 ZIP TREES

In recent work, Tarjan et al. [2021] introduce *zip trees* which can be seen as the functional equivalent of *skip lists* [Pugh 1990]. A zip tree is a binary search tree where every node also has a *rank*.

```
type ztree                                     alias key = int
  Leaf                                         alias rank = int
  Node(rank : rank, left : ztree, key : key, right : ztree)
```

We choose node ranks independently from a geometric distribution, where the rank of a node is non-negative integer k with probability $1/2^{k+1}$. Besides being a binary search tree for the keys, the tree is now also max heap-ordered with respect to the ranks with ties broken in favor of smaller keys. We define `is-higher-rank` as:




Fig. 3. Inserting a node with key 15 and rank 3 into a zip tree (with ranks shown as single digits in blue). Once the insertion point is found (as the right child of node 12), the tree at node 17 is *unzipped* along the key 15, and the resulting trees become the left- and right child of the inserted node. Deletion is the inverse where the children are *zipped* instead.

```
fip fun is-higher-rank( ^{r1,k1} : (rank,key), ^{r2,k2} : (rank,key) ) : bool
  (r1 > r2 || (r1 == r2 && k1 < k2))
```

Any parent node always *is-higher-rank* than its children. Since a zip tree is also a binary search tree, we can also see that the rank of a parent is always greater than the rank of its left-child, and greater than or equal to the rank of its right child. Interestingly, the shape of a zip tree is now fully determined by just the rank/key pairs in the tree and independent of the insertion order. See Figure 3 for an example of two valid zip trees. Intuitively we can see that given the geometric distribution of ranks, the shape of a tree naturally tends to be well balanced, with twice as many nodes at each lower rank. This means that the zip tree operations never need to do any explicit rebalancing, simplifying their implementation compared to usual balanced tree algorithms

The rank can be chosen independently at random, but in order to combine search and insertion, we can also derive the rank pseudo-randomly from the key. To insert an element into a zip tree, we first calculate the rank of the node. We can now traverse down until we find the fixed insertion point, as it is fully determined by the rank and key:

```
pub fun insert( t : ztree, k : key ) : ztree
  down( t, rank-of(k), k )
fun down( t : ztree, rank : rank, k : key ) : ztree
  match t
  Node(rnk,l,x,r) | is-higher-rank( (rnk,x), (rank,k) ) // go down while node is higher rank
    -> if (x < k) then Node(rnk, l, x, down(r,rank,k) )
        else Node(rnk, down(l,rank,k) , x, r)
  _ -> val (s,b) = unzip(t,k) in Node(rank,s,k,b)
```

Once we reach the insertion point where we are of higher rank than the current tree *t*, we *unzip* the tree *t* into two trees: one with elements smaller than *k* and one with bigger elements:

```
fun unzip( t : ztree, k : key ) : (ztree,ztree)
  match t
  Node(rnk,l,x,r) ->
    if (x < k) then val (s,b) = unzip(r,k) in (Node(rnk,l,x,s),b)
    elif (x > k) then val (s,b) = unzip(l,k) in (s,Node(rnk,b,x,r))
    else (l,r)
  Leaf -> (Leaf,Leaf)
```

Figure 3 illustrates inserting a node into a tree and the resulting unzip operation. Since the shape of a zip tree is always fixed by its rank/key pairs, *deletion* is the inverse of *insertion* which *zips* child trees back together. Like before, it is straightforward to formally prove that our specification of

insert maintains the expected properties of a zip tree.

6.1 Top-Down Zip Trees

To derive the top-down algorithm, we can again accumulate the smaller and bigger trees in constructor contexts. The unzip function becomes:

```
fun unzip( t : ztree, k : key, accl : ctx<ztree>, accr : ctx<ztree> ) : (ztree,ztree)
  match t
    Node(rnk,l,x,r) -> if (x < k) then unzip( r, k, accl ++ ctx Node(rnk,l,x,_), accr )
                       elif (x > k) then unzip( l, k, accl, accr ++ ctx Node(rnk,_,x,r) )
                       else (accl ++. l, accr ++. r)
    Leaf -> (accl ++. Leaf, accr ++. Leaf)
```

Unfortunately, this is not quite `fip` since in the case that the key is present in the unzipped tree, the `else` branch discards the `Node` on which we matched. However, we can avoid calling `unzip` in the first place if the key is present and derive an efficient `fip` version:

```
fip(1) fun insert-td( t : ztree, k : key ) : ztree
  down-td( t, rank-of(k), k, ctx _)

fip(1) fun down-td( t : ztree, rank : rank, k : key, acc : ctx<ztree> ) : ztree
  match t
    Node(rnk,l,x,r) | is-higher-rank( (rnk,x), (rank,k) )
      -> if (x < k) then down-td( r, rank, k, acc ++ ctx Node(rnk,l,x,_ ) )
          else down-td( l, rank, k, acc ++ ctx Node(rnk,_,x,r) )
    Node(_,_,x,_) | x == k -> acc ++. t
    _ -> val (s,b) = unzip-td( t, k, ctx _, ctx _ ) in acc ++. Node(rank,s,k,b)

fip fun unzip-td( t : ztree, k : key, accl : ctx<ztree>, accr : ctx<ztree> ) : (ztree,ztree)
  match t
    Node(rnk,l,x,r) -> if (x < k) then unzip-td( r, k, accl ++ ctx Node(rnk,l,x,_), accr )
                       else unzip-td( l, k, accl, accr ++ ctx Node(rnk,_,x,r) )
    Leaf -> (accl ++. Leaf, accr ++. Leaf)
```

It is straightforward to prove that the derived top-down algorithm is correct:

Theorem 8. (Correctness of top-down zip tree insertion)

$\text{down-td}(t,k,\text{acc}) \equiv \text{acc} ++. \text{insert}(t,k)$

The proof uses the following lemma for the correctness of `unzip-td`:

Lemma 1. (Correctness of top-down unzip)

$\text{unzip-td}(t,k,\text{accl},\text{accr}) \equiv \text{val } (s,b) = \text{unzip}(t,k) \text{ in } (\text{accl} ++. s, \text{accr} ++. b)$

6.2 Bottom-Up Zip Trees

For the bottom-up algorithm, we first define a zipper for the zip tree datatype to keep track of the path down the tree:

```
fip fun rebuild( z : zipper, t : ztree ) : ztree
  match z
    NodeR(rk,l,x,up)-> rebuild(up,Node(rk,l,x,t))
    NodeL(rk,up,x,r)-> rebuild(up,Node(rk,t,x,r))
    Done -> t
  type zipper
    NodeR(rk:rank, l:ztree, x:key, up:zipper)
    NodeL(rk:rank, up:zipper, x:key, r:ztree)
    Done
```

The `down` and `unzip` take the zipper(s) as an accumulating argument, where we again ensure we never unzip trees with the key present:

```
fip(1) fun insert-bu( t : ztree, k : key ) : ztree
  down-bu( t, rank-of(k), k, Done )

fip(1) fun down-bu( t : ztree, rank : rank, k : key, z : zipper ) : ztree
  match t
    Node(rnk,l,x,r) | is-higher-rank( (rnk,x), (rank,k) )
      -> if (x < k) then down-bu(r, rank, k, NodeR(rnk, l, x, z))
          else down-bu(l, rank, k, NodeL(rnk, z, x, r))
```




Fig. 4. Benchmarks on Ubuntu 22.04.2 (AMD 7950X 4.5Ghz) comparing the relative performance of C, ML, and Haskell against Koka for move-to-root (*mtr*), splay trees (*splay*), and zip trees (*zip*) for both top-down (*td*) and bottom-up (*bu*) variants. Each benchmark performs the same sequence of 10M pseudo-random insertions between 0 and 100 000 starting with an empty tree.

```

Node(_,_,x,_) | x == k -> rebuild(z, t)
_ -> val (s,b) = unzip-bu(t,k,Done,Done) in rebuild( z, Node(rank,s,k,b) )
fip fun unzip-bu( t : ztree, k : key, zs : zipper, zb : zipper ) : (ztree,ztree)
match t
  Node(rnk,l,x,r) ->
    if (x < k) then unzip-bu( r, k, NodeR(rnk,l,x,zs), zb)
    else unzip-bu( l, k, zs, NodeL(rnk,zb,x,r))
  Leaf -> (rebuild(zs,Leaf), rebuild(zb,Leaf))

```

We can optimize this a bit further: for the `down-bu` function, the zipper along the search path always just rebuilds the exact same path since no restructuring takes place, unlike the rebuilding for move-to-root or splay trees. It can be more efficient to use a constructor context for `down-bu` instead, as this can rebuild the tree in constant time.

For the optimized bottom-up version the correctness theorem is as before:

Theorem 9. (*Correctness of bottom-up zip tree insertion*)

`down-bu(t,k,acc) ≡ acc ++. insert(t,k) with`

`unzip-bu(t,k,zs,zb) ≡ val (s,b) = unzip(t,k) in (rebuild(zs,s), rebuild(zb,b))`

As with move-to-root and splay trees, we can again prove the correctness of imperative zip insertion. But this time we can even go further: In Appendix C in the tech report, we present a new imperative insertion algorithm which derives directly from the functional version and is simpler, yet as-efficient as the imperative algorithm by Tarjan et al. [2021].

7 BENCHMARKS

Figure 4 shows benchmark results for the various derived algorithms in this paper. We compare Koka against the best known iterative C implementations. For bottom-up algorithms, we also benchmark ML and Haskell implementations that are direct translations of the bottom-up Koka versions. We ran the benchmarks on Ubuntu 22.04.2 using an AMD 7950X at 4.5Ghz. We used Koka v2.4.2 (`-o2`), the C implementations were compiled with Clang 14 (`-o3 -DNDEBUG`), ML with OCaml 4.13.1 (`ocamltop -o2`), and Haskell with GHC 8.8.4 (`-o2`). Each bar represents the median performance over 10 runs, with error bars for the standard error. Each benchmark performs 10M insertions starting with an empty tree, using a pseudo random sequence of keys between 0 and 100 000. Initially the tree is populated quickly up to 100 000 elements followed by many insertions

where the element already exists. We tested all top-down (*-td*) and bottom-up (*-bu*) versions of move-to-root tree (*mtr*), splay trees (*splay*), and zip trees (*zip*). Figure 4 also includes tests for red-black trees (*rb*) but we disregard those for the moment.

If we look at the performance relative to the median performance of Koka in Figure 4 we see that our purely functional *fip* derived versions always outperform C for move-to-root, splay, and zip-trees! How is that even possible? The Koka code in particular must perform more operations:

- Koka has automatic memory management, and thus everything is reference counted. The generated code also includes branches to handle potential thread shared structures (which requires atomic reference count operations).
- Koka uses arbitrary precision integers (*int*) for keys and all comparisons and arithmetic operations include branches for the case where big integer arithmetic is required.
- Context composition and application are reference counted to handle sharing, and always check for empty contexts. In the C code empty context checks are unnecessary due to stack allocation.
- Koka reuses memory when possible, but the trees can always be used *persistently* as well, and insertion can also handle shared trees where the spine to the insertion point is copied.

One factor why Koka still outperforms C is that Koka is tightly integrated with the optimized *mimalloc* allocator [Leijen et al. 2019]. To gain better insight into what the actual overhead of the above features is in our functional code, we also include “equalized” C: here we link the C programs with *mimalloc* as well (overriding *malloc* and *free*), and we include an unused header word in the top-down algorithms to ensure an equal amount of memory is allocated.² This is the third bar in Figure 4. Even compared to equalized C our functional versions still perform remarkably well, being at most 15% slower for top-down move-to-root trees, and only 6% slower for top-down zip-trees. This is surprising, given the additional safety guarantees Koka provides. Many of these checks are *cache-local* and use just few instructions in the fast path (e.g. *is-unique*). We conjecture that on modern hardware small fast-path branches with cache-local accesses can be quite cheap – due to the speculation with many parallel compute units the actual performance bottlenecks may be somewhere else, such as a dependency on an uncached memory read.

Even with equalized C, our functional versions are still substantially faster on the bottom-up move-to-root and splay trees. This is due to the difference in implementations: in our derived functional versions we use zippers which are compiled essentially to use in-place pointer-reversal at runtime. The C implementations, in contrast, are using parent pointers instead which is the usual way of traversing back up for the bottom-up algorithms. However, for move-to-root and splay trees the constant restructuring is now more expensive since we need to also adjust parent pointers for each rotation. This cost is much less pronounced for zip trees for which considerably less restructuring takes place, and so the performance difference is correspondingly smaller. As an experiment, we also implemented a pointer-reversal version of Allen and Munro’s move-to-root bottom-up algorithm using the lowest pointer bit to distinguish left- from right paths. In that case, the equalized C code performs about 14% better than our functional version.

The top-down zip tree algorithm in C uses Tarjan et al.’s algorithm. We also tested this with our derived algorithm, and the simpler version that does not have the inner *repeat-until* iterations (and may thus perform extra pointer assignments as shown in Appendix C). For our benchmark though, we could not measure any significant differences in execution times between these versions.

Red-Black Trees. Figure 4 also contains results for bottom-up [Guibas and Sedgewick 1978; Lorenzen and Leijen 2022; Okasaki 1999a] and top-down [Tarjan 1985; Weiss 2013] red-black tree algorithms. It is beyond the scope of this paper to discuss those in detail but we can apply the

²This is not required for the bottom-up algorithms in C since these have parent pointers which balances out against Koka’s header words (which uses pointer reversal through zippers and needs no parent pointers).

same techniques that we have shown in this paper to implement the bottom-up version using defunctionalized CPS and zippers, and the top-down version with constructor contexts. The top-down C version is based on the GNU library tree search implementation which encodes the node color in the least significant pointer bit [Schmidt 1997], while the bottom-up one implements the algorithm described by Cormen et al. [2022]. The relative performance of Koka versus (equalized) C is still good, but less impressive as for the other data structures: about 25% slower for the top-down algorithm and almost 50% slower for the bottom-up version. This shows that there is still room for further improvements in our compilation techniques.

Each variant performs poorly for different reasons though. We believe the functional version of bottom-up red-black trees is slower because the C versions can use early bailout: on the way back up as soon as a parent is no longer red, the C version can immediately return the root pointer. For the functional version though we need to unwind the zipper completely to reconstruct the tree. There seems no obvious way to implement such optimization on the functional side: we would need some concept of parent pointers to achieve similar behaviour. For the top-down version the reason for the poorer performance is less clear, but we believe it is due to the need to keep track of extra context. Top-down red-black tree rebalancing requires access to the parent and grand-parent of the current node for its rebalancing operations. In C we can just keep two extra pointers around on the traversal down. In the functional version though we need two derivative node constructors for the parent and grandparent, together with the accumulating constructor context – moving the grand-parent into the constructor context on each iteration. We imagine that a potential path to improving this situation is to allow a limited form of pattern matching on constructor contexts.

8 RELATED WORK

We discuss related work of the studied algorithms in the main text. Here, we want to present an overview of the work related to the employed techniques.

Data structures with a hole. Zippers [Huet 1997] are the canonical functional representation of data structures with a hole. They can be defined generically [Hinze et al. 2004; McBride 2001 2008], but also arise syntactically as the defunctionalization of the closures generated by a CPS-conversion [Danvy and Nielsen 2001]. While they have long been known to be the functional equivalent of backpointers [Huet 2003; Munch-Maccagnoni and Douence 2019], only recently has this insight been exploited to actually compile them to pointer reversing code [Lorenzen et al. 2023b].

In contrast, constructor contexts, as studied in this work, have received far less attention. One reason for this may be that previous implementations required type systems to ensure safety. Minamide [1998] describes a linear type system for efficient one-hole contexts, while destination passing style [Bour et al. 2021; Pottier and Protzenko 2013] requires linear or ownership types. Huet [2003] also discusses top-down structures with a hole Ω , but he does not make an explicit connection to top-down algorithms or present an efficient implementation.

Some top-down algorithms can also be expressed using either laziness [Wadler 1984] or tail recursion modulo cons (TRMC) [Bour et al. 2021; Friedman and Wise 1975; Leijen and Lorenzen 2023; Risch 1973]. However, both techniques require the programmer to provide an expression up-front which determines the value eventually filling the hole. This makes it more cumbersome and sometimes impossible to express top-down algorithms with these techniques. Laziness additionally carries a performance overhead due to the creation of intermediate thunks. Conversely, TRMC can be implemented manually with first-class constructor contexts: Leijen and Lorenzen [2023] introduce the *context transformation*, which generalises Danvy and Nielsen’s [2001] approach to constructor contexts.

Compilation of functional programs. A crucial step of our compilation is to *reuse* [Lorenzen and Leijen 2022; Schulte and Grieskamp 1992; Ullrich and de Moura 2019] old heap cells for new ones. This can be performed automatically in languages with precise reference counting [Reinking, Xie et al. 2021; Ullrich and de Moura 2019], but could also be manually implemented in languages with uniqueness types [Barendsen and Smetzers 1996]. However, in order to achieve a fully in-place algorithm, we also need to be sure that certain values (such as tuples) are not allocated on the heap. Lorenzen et al. [2023b] propose a calculus for such functions which ensures that the functions presented here do not have spurious allocations.

In this work, we study compilation as a refinement [Appel 2016] which allows us to connect the functional implementation to published imperative code. Modulo exact choice of variable names and helper functions, it is possible to *compile* functional code directly to published imperative code. Hofmann [2000] first proposed such a scheme and Gudjónsson and Winsborough [1999] presents an optimization to avoid updating the hole of the context in cases where it already contains the right value, just as in the published implementation of zip tree insertion.

Verification of imperative algorithms. Insertion and deletion algorithms for binary search trees have been verified countless times: There is a large literature on functional implementations [Nipkow et al. 2021 2020] as well as destructive implementations [Armborst and Huisman 2021; Pek et al. 2014; Stefanescu et al. 2016 2016; Zhan 2018]. However, these algorithms are typically based on *recursive* code and thus do not deal with the issues discussed in this paper. Surprisingly, there seems to be far less literature on verifying idiomatic, imperative code as it appears in algorithm papers. Schellhorn et al. [2022] and Dross and Moy [2017] formalize the text-book insertion and deletion of red-black trees, but due to the use of inline invariants their code does not resemble the original implementation. Lammich [2020] formalizes an array-based implementation of pattern-defeating quicksort in the Boost C++ library. Enea et al. [2015] prove insertion algorithms for AVL trees and red-black trees in C correct by deriving a representation for the already-traversed segment. They do not consider a functional version and thus have to perform a proof search.

Formalizing constructor contexts. Following Charguéraud [2016], we define an inductive representation of one-hole data structures. In the work of Enea et al. [2015], these *segments* also hold additional invariants, but this is not necessary if one only wants to relate the segments to their functional counterpart. Cao et al. [2019] formalize an idiomatic, non-balancing insertion into binary trees and point out that a constructor context can also be represented in separation logic using a magic wand, thereby implementing constructor contexts directly as their interface.

Tuerk [2010] demonstrates a method for proving the correctness of while-loops that recurse on an argument by using simple pre- and post-conditions; this may be powerful enough to prove bottom-up algorithms as well as those top-down algorithms that arise from a functional version via TRMC.

9 CONCLUSION

This work bridges the gap between imperative algorithms and purely functional programs. The key techniques to guarantee performant functional implementations—deriving tail recursive fip functions from their directly recursive counterparts—are in no way restricted to binary search trees. We fully expect them to be widely applicable to other algorithms and data structures. Furthermore, this will enable us to adopt a wide variety of techniques designed specifically for functional languages—ranging from program synthesis [Albarghouthi et al. 2013; Polikarpova et al. 2016] to automatic amortized complexity analysis [Leutgeb et al. 2022; Schoenmakers 1993]—in a novel setting, where they could not be applied easily until now.

ACKNOWLEDGEMENTS

This work was carried out while Anton Lorenzen was a Research Intern at Microsoft Research. Sam Lindley was supported by UKRI Future Leaders Fellowship MR/T043830/1 (EHOP).

REFERENCES

- Adams. 1993. Functional Pearls Efficient Sets—a Balancing Act. *Journal of Functional Programming* 3 (4). Cambridge University Press: 553–561. doi:<https://doi.org/10.1017/S095679680000885>.
- Albarghouthi, Gulwani, and Kincaid. 2013. Recursive Program Synthesis. In *Computer Aided Verification: 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13–19, 2013. Proceedings 25*, 934–950. Springer. doi:https://doi.org/10.1007/978-3-642-39799-8_67.
- Allen, and Munro. 1978. Self-Organizing Binary Search Trees. *Journal of the ACM (JACM)* 25 (4). ACM New York, NY, USA: 526–535. doi:<https://doi.org/10.1145/322092.322094>.
- Appel. 2016. Modular Verification for Computer Security. In *2016 IEEE 29th Computer Security Foundations Symposium (CSF)*, 1–8. IEEE. doi:<https://doi.org/10.1109/CSF.2016.8>.
- Appel. 2018. Software Foundations Volume 3: Verified Functional Algorithms. Electronic textbook.
- Arnborst, and Huisman. 2021. Permission-Based Verification of Red-Black Trees and Their Merging. In *2021 IEEE/ACM 9th International Conference on Formal Methods in Software Engineering (FormalISE)*, 111–123. IEEE. doi:<https://doi.org/10.1109/FormaliSE52586.2021.00017>.
- Aspinall, Hofmann, and Konečný. 2008. A Type System with Usage Aspects. *Journal of Functional Programming* 18 (2). Cambridge University Press: 141–178. doi:<https://doi.org/10.1017/S0956796807006399>.
- Barendsen, and Smetsers. 1996. Uniqueness Typing for Functional Languages with Graph Rewriting Semantics. *Mathematical Structures in Computer Science* 6 (6). Cambridge University Press: 579–612. doi:<https://doi.org/10.1017/S0960129500070109>.
- Blelloch, Ferizovic, and Sun. 2016. Just Join for Parallel Ordered Sets. In *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures*, 253–264. doi:<https://doi.org/10.1145/2935764.2935768>.
- Bour, Clément, and Scherer. Apr. 2021. Tail Modulo Cons. *Journées Francophones Des Langages Applicatifs (JFLA)*, April. Saint Médard d’Excideuil, France. <https://hal.inria.fr/hal-03146495/document>. hal-03146495.
- Cao, Wang, Hobor, and Appel. 2019. Proof Pearl: Magic Wand as Frame. *arXiv Preprint arXiv:1909.08789*.
- Charguéraud. 2016. Higher-Order Representation Predicates in Separation Logic. In *Proceedings of the 5th ACM SIGPLAN Conference on Certified Programs and Proofs*, 3–14. doi:<https://doi.org/10.1145/2854065.2854068>.
- Clark, and Tärnlund. 1977. A First Order Theory of Data and Programs. In *IFIP Congress*, 939–944.
- Cormen, Leiserson, Rivest, and Stein. 2022. *Introduction to Algorithms*. MIT press.
- Danvy, and Nielsen. 2001. Defunctionalization at Work. In *Proceedings of the 3rd ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming*, 162–174. doi:<https://doi.org/10.1145/773184.773202>.
- Dross, and Moy. 2017. Auto-Active Proof of Red-Black Trees in SPARK. In *NASA Formal Methods: 9th International Symposium, NFM 2017, Moffett Field, CA, USA, May 16–18, 2017, Proceedings 9*, 68–83. Springer. doi:https://doi.org/10.1007/978-3-319-57288-8_5.
- Enea, Sighireanu, and Wu. 2015. On Automated Lemma Generation for Separation Logic with Inductive Definitions. In *Automated Technology for Verification and Analysis: 13th International Symposium, ATVA 2015, Shanghai, China, October 12–15, 2015, Proceedings 13*, 80–96. Springer. doi:https://doi.org/10.1007/978-3-319-24953-7_7.
- Erbsen, Gruetter, Choi, Wood, and Chlipala. 2021. Integration Verification across Software and Hardware for a Simple Embedded System. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, 604–619. doi:<https://doi.org/10.1145/3410295>.
- Friedman, and Wise. Dec. 1975. *Unwinding Stylized Recursion into Iterations*. 19. Bloomington, Indiana. <https://legacy.cs.indiana.edu/ftp/techreports/TR19.pdf>.
- Frumin, Gondelman, and Krebbers. 2019. Semi-Automated Reasoning About Non-Determinism in C Expressions. In *ESOP*, 60–87. doi:https://doi.org/10.1007/978-3-030-17184-1_3.
- Gudjónsson, and Winsborough. 1999. Compile-Time Memory Reuse in Logic Programming Languages through Update in Place. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 21 (3). ACM New York, NY, USA: 430–501. doi:<https://doi.org/10.1145/319301.319309>.
- Guibas, and Sedgewick. 1978. A Dichromatic Framework for Balanced Trees. In *19th Annual Symposium on Foundations of Computer Science (sfcs 1978)*, 8–21. IEEE. doi:<https://doi.org/10.1109/SFCS.1978.3>.
- Hinze, Jeuring, and Löh. 2004. Type-Indexed Data Types. *Science of Computer Programming* 51 (1–2). Elsevier: 117–151. doi:https://doi.org/10.1007/3-540-45442-X_10.
- Hofmann. 2000. In-Place Update with Linear Types or How to Compile Functional Programs into Malloc-Free C. *Preprint*. Citeseer.
- Huet. 1997. The Zipper. *Journal of Functional Programming* 7 (5): 549–554. doi:<https://doi.org/10.1017/S0956796897002864>.

- Huet. 2003. Linear Contexts, Sharing Functors: Techniques for Symbolic Computation. *Thirty Five Years of Automating Mathematics*. Springer, 49–69. doi:https://doi.org/10.1007/978-94-017-0253-9_4.
- Hughes. 1986. A Novel Representation of Lists and Its Application to the Function “reverse.” *Information Processing Letters* 22 (3). Elsevier: 141–144. doi:[https://doi.org/10.1016/0020-0190\(86\)90059-1](https://doi.org/10.1016/0020-0190(86)90059-1).
- Jung, Krebbers, Jourdan, Bizjak, Birkedal, and Dreyer. 2018. Iris from the Ground up: A Modular Foundation for Higher-Order Concurrent Separation Logic. *Journal of Functional Programming* 28. Cambridge University Press: e20. doi:<https://doi.org/10.1017/S0956796818000151>.
- Lammich. 2020. Efficient Verified Implementation of Introsort and Pdqsort. In *Automated Reasoning: 10th International Joint Conference, IJCAR 2020, Paris, France, July 1–4, 2020, Proceedings, Part II 10*, 307–323. Springer. doi:https://doi.org/10.1007/978-3-030-51054-1_18.
- Leijen. 2014. Koka: Programming with Row Polymorphic Effect Types. In *MSFP’14, 5th Workshop on Mathematically Structured Functional Programming*. doi:<https://doi.org/10.4204/EPTCS.153.8>.
- Leijen. 2021. The Koka Language. <https://koka-lang.github.io>.
- Leijen, and Lorenzen. Jul. 2022. *Tail Recursion Modulo Context – An Equational Approach*. MSR-TR-2022-18. Microsoft Research.
- Leijen, and Lorenzen. Jan. 2023. Tail Recursion Modulo Context: An Equational Approach. *Proc. ACM Program. Lang.* 7 (POPL). doi:<https://doi.org/10.1145/3571233>. See also [Leijen and Lorenzen 2022].
- Leijen, Zorn, and de Moura. 2019. Mimalloc: Free List Sharding in Action. *Programming Languages and Systems, LNCS*, 11893. Springer International Publishing. doi:https://doi.org/10.1007/978-3-030-34175-6_13. APLAS’19.
- Leutgeb, Moser, and Zuleger. 2022. Automated Expected Amortised Cost Analysis of Probabilistic Data Structures. In *International Conference on Computer Aided Verification*, 70–91. Springer. doi:https://doi.org/10.1007/978-3-031-13188-2_4.
- Lorenzen, and Leijen. Sep. 2022. Reference Counting with Frame Limited Reuse. In *Proceedings of the 27th ACM SIGPLAN International Conference on Functional Programming (ICFP’22)*. ICFP’22. Ljubljana, Slovenia. doi:<https://doi.org/10.1145/3547634>.
- Lorenzen, Leijen, and Swierstra. May 2023. *FP²: Fully in-Place Functional Programming*. MSR-TR-2023-19. Microsoft Research.
- Lorenzen, Leijen, and Swierstra. Sep. 2023. *FP²: Fully in-Place Functional Programming*. In *Proceedings of the 28th ACM SIGPLAN International Conference on Functional Programming (ICFP’2023)*. ICFP’23. Seattle, USA. Under submission. See [Lorenzen et al. 2023a] for the extended technical report.
- Lorenzen, Leijen, Swierstra, and Lindley. Mar. 2024. The Functional Essence of Imperative Binary Search Trees (Artifact). Zenodo. doi:<https://doi.org/10.5281/zenodo.10790231>. Artifact for PLDI’24.
- Lucas. May 2004. A Direct Algorithm for Restricted Rotation Distance. *Inf. Process. Lett.* 90 (3). Elsevier North-Holland, Inc.: 129–134. doi:<https://doi.org/10.1016/j.ipl.2004.02.001>.
- McBride. 2001. The Derivative of a Regular Type Is Its Type of One-Hole Contexts. <http://strictlypositive.org/diff.pdf>. (Extended Abstract).
- McBride. 2008. Clowns to the Left of Me, Jokers to the Right (pearl) Dissecting Data Structures. In *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 287–295. doi:<https://doi.org/10.1145/1328897.1328474>.
- Minamide. 1998. A Functional Representation of Data Structures with a Hole. In *Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 75–84. POPL ’98. San Diego, California, USA. doi:<https://doi.org/10.1145/268946.268953>.
- Mulder, Czajka, and Krebbers. 2023. Beyond Backtracking: Connections in Fine-Grained Concurrent Separation Logic. *Proceedings of the ACM on Programming Languages* 7 (PLDI). ACM New York, NY, USA: 1340–1364. doi:<https://doi.org/10.1145/3591275>.
- Mulder, Krebbers, and Geuvers. 2022. Diaframe: Automated Verification of Fine-Grained Concurrent Programs in Iris. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, 809–824. doi:<https://doi.org/10.1145/3519939.3523432>.
- Munch-Maccagnoni, and Douence. 2019. Efficient Deconstruction with Typed Pointer Reversal. In *ML 2019-Workshop*, 1–8.
- Nipkow, Blanchette, Eberl, Gómez-Londoño, Lammich, Sternagel, Wimmer, and Zhan. 2021. Functional Algorithms, Verified.
- Nipkow, Eberl, and Haslbeck. 2020. Verified Textbook Algorithms: A Biased Survey. In *Automated Technology for Verification and Analysis: 18th International Symposium, ATVA 2020, Hanoi, Vietnam, October 19–23, 2020, Proceedings 18*, 25–53. Springer. doi:https://doi.org/10.1007/978-3-030-59152-6_2.
- Okasaki. 1999. Red-Black Trees in a Functional Setting. *Journal of Functional Programming* 9 (4). Cambridge University Press: 471–477. doi:<https://doi.org/10.1017/S0956796899003494>.
- Okasaki. Jun. 1999. *Purely Functional Data Structures*. Columbia University, New York.
- Pek, Qiu, and Madhusudan. 2014. Natural Proofs for Data Structure Manipulation in C Using Separation Logic. *ACM SIGPLAN Notices* 49 (6). ACM New York, NY, USA: 440–451. doi:<https://doi.org/10.1145/2666356.2594325>.

- Pit-Claudel, Philipoom, Jamner, Erbsen, and Chlipala. 2022. Relational Compilation for Performance-Critical Applications: Extensible Proof-Producing Translation of Functional Models into Low-Level Code. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, 918–933. doi:<https://doi.org/10.1145/3519939.3523706>.
- Polikarpova, Kuraj, and Solar-Lezama. 2016. Program Synthesis from Polymorphic Refinement Types. *ACM SIGPLAN Notices* 51 (6). ACM New York, NY, USA: 522–538. doi:<https://doi.org/10.1145/2908080.2908093>.
- Pottier, and Protzenko. 2013. Programming with Permissions in Mezzo. In *Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming*, 173–184. ICFP '13. ACM, Boston, Massachusetts, USA. doi:<https://doi.org/10.1145/2500365.2500598>.
- Pugh. 1990. Skip Lists: A Probabilistic Alternative to Balanced Trees. *Communications of the ACM* 33 (6). ACM New York, NY, USA: 668–676.
- Reinking, Xie, de Moura, and Leijen. 2021. Perceus: Garbage Free Reference Counting with Reuse. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, 96–111. PLDI 2021. ACK, New York, NY, USA. doi:<https://doi.org/10.1145/3453483.3454032>.
- Reynolds. 1972. Definitional Interpreters for Higher-Order Programming Languages. In *Proceedings of the ACM Annual Conference - Volume 2*, 717–740. ACM, Boston, Massachusetts, USA. doi:<https://doi.org/10.1145/800194.805852>.
- Reynolds. 2002. Separation Logic: A Logic for Shared Mutable Data Structures. In *Proceedings 17th Annual IEEE Symposium on Logic in Computer Science*, 55–74. IEEE. doi:<https://doi.org/10.1109/LICS.2002.1029817>.
- Risch. Nov. 1973. REMREC - A Program for Automatic Recursion Removal. Inst. för Informationsbehandling, Uppsala Universitet. <https://user.it.uu.se/~torer/publ/remrec.pdf>.
- Schellhorn, Bodenmüller, Bitterlich, and Reif. 2022. Separating Separation Logic—modular Verification of Red-Black Trees. In *Working Conference on Verified Software: Theories, Tools, and Experiments*, 129–147. Springer. doi:https://doi.org/10.1007/978-3-031-25803-9_8.
- Schmidt. 1997. The GNU C Library, tsearch. <https://github.com/lattera/glibc/blob/master/misc/tsearch.c>.
- Schoenmakers. 1993. A Systematic Analysis of Splaying. *Information Processing Letters* 45 (1). Elsevier: 41–50. doi:[https://doi.org/10.1016/0020-0190\(93\)90249-9](https://doi.org/10.1016/0020-0190(93)90249-9).
- Schorr, and Waite. 1967. An Efficient Machine-Independent Procedure for Garbage Collection in Various List Structures. *Communications of the ACM* 10 (8). ACM New York, NY, USA: 501–506. doi:<https://doi.org/10.1145/363534.363554>.
- Schulte, and Grieskamp. 1992. Generating Efficient Portable Code for a Strict Applicative Language. In *Declarative Programming, Sasbachwalden 1991*, 239–252. Springer. doi:https://doi.org/10.1007/978-1-4471-3794-8_16.
- Sleator, and Tarjan. 1985. Self-Adjusting Binary Search Trees. *Journal of the ACM (JACM)* 32 (3). ACM New York, NY, USA: 652–686. doi:<https://doi.org/10.1145/3828.3835>.
- Stefanescu, Park, Yuwen, Li, and Roşu. 2016. Semantics-Based Program Verifiers for All Languages. *ACM SIGPLAN Notices* 51 (10). ACM New York, NY, USA: 74–91. doi:<https://doi.org/10.1145/3022671.2984027>.
- Stephenson. 1980. A Method for Constructing Binary Search Trees by Making Insertions at the Root. *International Journal of Computer & Information Sciences* 9. Springer: 15–29. doi:<https://doi.org/10.1007/BF00995807>.
- Tarjan. May 1985. *Efficient Top-Down Updating of Red-Black Trees*. TR-006-85. Princeton University. <https://www.cs.princeton.edu/research/techreps/TR-006-85>.
- Tarjan, Levy, and Timmel. 2021. Zip Trees. *ACM Transactions on Algorithms (TALG)* 17 (4). ACM New York, NY: 1–12. doi:<https://doi.org/10.1145/3476830>.
- The Coq Development Team. Oct. 2017. The Coq Proof Assistant, Version 8.7.0. Zenodo. doi:<https://doi.org/10.5281/zenodo.1028037>.
- The Iris Team. 2022. *The Iris 4.0 Reference*. <https://plv.mpi-sws.org/iris>.
- Tuerk. 2010. Local Reasoning about While-Loops. *VSTTE 2010*: 29.
- Ullrich, and de Moura. Sep. 2019. Counting Immutable Beans – Reference Counting Optimized for Purely Functional Programming. In *Proceedings of the 31st Symposium on Implementation and Application of Functional Languages (IFL'19)*. Singapore. doi:<https://doi.org/10.1145/3412932.3412935>.
- Wadler. 1984. Listlessness Is Better than Laziness: Lazy Evaluation and Garbage Collection at Compile-Time. In *Proceedings of the 1984 ACM Symposium on LISP and Functional Programming*, 45–52. doi:<https://doi.org/10.1145/800055.802020>.
- Weiss. 2013. *Data Structures and Algorithm Analysis in C++ (Fourth Edition)*. Addison-Wesley.
- Zhan. 2018. Efficient Verification of Imperative Programs Using auto2. In *Tools and Algorithms for the Construction and Analysis of Systems: 24th International Conference, TACAS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings, Part I 24*, 23–40. Springer. doi:https://doi.org/10.1007/978-3-319-89960-2_2.